

Network Working Group
Request for Comments #451
NIC #14135

M. A. Padjipsky
MIT-Multics
February 22, 1973

Tentative Proposal for a Unified User Level Protocol

Now that proposals for expansions to the Telnet Protocol are in vogue again (RFC's 426 and 435, for example), I'd like to promote some discussion of a particular favorite of my own. Please note that this is presented as a tentative proposal: it's an attempt to consider the desirability of a new approach, not a rigorous specification. To begin somewhat obliquely, for some time I've felt that we (the NWG) have fallen into a trap in regard to the Initial Connection Protocol. The point is that even though the ICP gives us the ability to define a "family" of ICPlets by varying the contact socket, there's no compelling reason why we should do so. That we have done so in the FTP and RJEP I view as unfortunate--but also undesirable and unnecessary.

To take the "undesirable" aspects first, consider the following: If we continue to define a new contact socket for every new "user level" protocol we come up with, we'll continue to need another new mechanism (process, procedure, or patch) to respond to requests for connection for each new protocol. By Occam's Razor (or the principle of economy of mechanism, if you prefer), this is a bad thing. Irrespective of the relative difficulty of implementing such mechanisms on the various Hosts, to implement them at all leads to a kind of conceptual clutter. Further, a different kind of confusion is introduced by the notion which some of our number seem to be entertaining, that the "later" user level protocols such as FTP are somehow still another level of abstraction up from Telnet. So it seems to me that we could spare ourselves a lot of bother, both practical and theoretical, if we could avoid spawning contact sockets needlessly.

Turning to the "unnecessary" aspects, I think that even if the case against the current approach isn't completely convincing the case for a particular alternative might be. So to show that the multiple contact socket ICP is unnecessary, I'll try to show that what I call the "unified user level protocol" (UULP) is better. The first thing to notice is that all the "later" protocols "speak Telnet". This is sensible: Telnet works, by and large. Why not make use of it? Right. But why not make even more use of it? In view of the fact that FTP, RJEP, and even the initiating part of the Network Graphics Protocol, are really just ways of letting a user say to a Server "I don't know what you call it on your system, but please perform the whatever function (push or pull a file, start or stop a batch job, funnel some of my output through the Network Virtual Graphics Terminal module) for me now," why not simply put hooks in Telnet to indicate that a Network Generic Function is wanted instead of a Host-specific one at a given point in time? Then everybody can come in through Telnet in ways that are already known (and usually debugged and optimized) and fan out to other services through a single mechanism, where that single mechanism can be whatever is most

own that the
acterized by
data is sent or
1 after which
of dedicated
in very ineffi-
acket of some
munication for
transfer data in
length streams
nts correspond
segment could

Computer ca-
pabilities infor-
in the research
achieved in this

individual per-
ne network such
nically allocated
et. Each packet
e network and
and is delivered
illy a fraction of
dual user traffic.

for that user is
ould also lead to
even if the alloca-
tion is not ideal.

ability to provide
partants. Full com-
can communica-
ce and adminis-
tration and par-
ameters or con-
reliability are
of interconnected
vironment. This
because lower er-
r, in turn, im-
rease, the delay
propagation delay
xtends the original

for networks
domain of broad-
cast in this area
provide computer
computer commu-
radio offers a high
channel, particu-
bers of users

bursty traffic. During the early 1970's, the ALOHA project at the University of Hawaii demonstrated the feasibility of using packet broadcasting in a single-hop system (see reference [8], [9]). The Hawaii work led to the development of a multi-hop multiple access packet radio network (PRNET) under the sponsorship of the Advanced Research Projects Agency (ARPA). The PRNET is a fundamental network extension of the basic ALOHA system and broadens the realm of packet communications to permit mobile applications over a wide geographic area. The use of broadcast radio technology for local distribution of information can also provide a degree of flexibility in rapid deployment and reconfiguration not currently possible with most fixed plant installations. Although the original impetus for packet radio development was and still is largely based on tactical military computer communication requirements [10], the basic concept is applicable to an extremely wide range of new and innovative computer communication applications never before possible in any practical way.

In addition to the strong ARPANET and ALOHA system influences, three technical developments in the early 1970's were largely responsible for the evolution of packet switching to the radio environment. The first was the microprocessor and associated memory technology which made it possible to incorporate computer processing at each packet radio network node in a form that was compatible with mobile usage and portable operation. The second was the reduction to practice of surface acoustic wave (SAW) technology which can perform matched filtering (to receive wide-band radio signals) on a very small substrate of quartz or similar piezo-electric material. The third development was purely conceptual and involved an awareness within the computer and communications communities of the importance of "protocols" in the development of network management strategies. It was upon those three pillars that the technical approach to the PRNET was founded.

Packet radio network technology will be essential for military and other governmental needs as terminals and computer systems become pervasive throughout essentially all aspects of their operations. Initially, the needs for radio-based computer communications are expected to be prevalent in training, on or near the battlefield, and for crisis situations. The first operational systems, even if of limited availability, are most likely to be deployed for use in one of these areas where the higher relative cost of providing the advanced capability can be tolerated. Within the civilian sector, there is also a strong need for terminal access to information in the mobile environment, but the cost of service to the user (e.g., personal terminal, tariffs) will dictate when such capabilities should be publicly provided. We expect to see a considerable increase in the usage of civilian terminals and microcomputers "on the move" during the early 1980's but, in contrast to the military environment, these applications are expected to involve relatively simple equipment, reduced capabilities and lower costs.

All users in a packet radio network are assumed to share a common radio channel, access to which is controlled by microprocessors in the packet radios. In contrast to a CB radio channel in which contention for the channel is directly controlled by the users (who at best can do a poor job of scheduling the channel), the packet radio system decouples direct access to the channel from user requests for channel access. Within a fraction of a second, the microprocessors can dynamically schedule and control the channel to minimize or avoid conflicts (overlapping transmissions) particularly when

the transmissions are very short. The use of computer control for channel access can lead to very efficient system operation relative to other more conventional manual methods of access control [11].

In recent years, the subject of efficient spectrum utilization has received increasing attention. A special issue of *IEEE Transactions on Electromagnetic Compatibility* on spectrum management [12] addresses this topic in considerable detail. This subject is not addressed further in this paper, other than to note that because of its capability for dynamic allocation of the spectrum, packet radio is a particularly good choice to obtain efficient utilization for bursty traffic. The ability to achieve effective usage of the spectrum will be a central factor along with cost in determining the ultimate viability of radio based networks for local distribution of information.

In this paper, we discuss the basic concepts of packet radio and present the recent technology and system advances. In Section II, we indicate various capabilities and services which a packet radio network might provide. In Section III, we consider the problem of signaling over a ground radio channel with all its attendant environmental factors. In Section IV, we discuss the basic operation of a packet radio network with underlying emphasis on the network elements and system protocols. In Section V, we discuss several advanced system capabilities for operation and control of the packet radio system. In Section VI, we separate out a subject of particular interest, namely spread spectrum transmission in the network environment. Although a packet radio system need not employ spread spectrum, there are several noteworthy attributes arising from its use. The experimental packet radio network (PRNET) under development by the ARPA is discussed in Section VII. The final section contains conclusions.

II. CAPABILITIES AND SERVICES OF A PR NETWORK

A primary objective of a packet radio network is to support real-time interactive communications between computer resources (hosts) connected to the network and user terminals (e.g., terminal-host, host-host, and terminal-terminal interactions). In order to satisfy this objective, the network should provide certain basic capabilities and services which can be grouped roughly into two categories: those which are always or automatically provided by the network and those which a user may select based on his application. The former category includes such capabilities as network transparency, area coverage/connectivity, mobile operation, internetting, coexistence, throughput with low delay, and rapid deployment. The last category includes error control options, routing options, addressing options, and services for various tactical applications.

We identify here a few of these basic packet radio network services and capabilities. While the list is not intended to be exhaustive, those items on it are all major factors of interest.

We assume that computer resources (hosts) need to be connected with each other and with individual users who might access data bases, manipulate files, run programs or write and execute programs to run on remote hosts. The packet radio network merely provides a high throughput, low delay means of interconnection for the (potentially mobile) community of users. Many of these operations will be interactive, with a computer response to a remote user entry being desired in real-time. Although the primary objective of the net is to provide service to computer communication traffic, other types of service, such as might be required for real-time

speech, can be accommodated along with the capability for end-to-end security based on packet encryption techniques.

A. Transparency

The basic internal operation of the network should be transparent to the user. We use this term to mean that all user data presented to the net should be delivered to the destination without modification of the information content in any way. Only the data to be delivered, and the necessary control and addressing information should be required of the user as input. All other aspects of routing, reliable delivery, protocols and network operation should be handled by the network itself. Only in the case of communication difficulties should the user or user process be advised of internal network status. Transparency is desirable in order to allow the network to regulate and optimize its internal flow of traffic in a global manner without unnecessary constraints applied by users. The users, in turn, need not be concerned with the activities taking place in the net or their effect on network operations, but need only specify the services desired.

B. Area Coverage and Connectivity

Area coverage with full connectivity should be provided. For ground mobile radio, network diameters on the order of 100 miles are appropriate, but the system architecture should allow the geographic area of coverage to be expanded at the expense of increased end-to-end delay across the network. All valid traffic originators within the net must be provided connectivity with all other valid receivers subject only to the overall reliability and performance of the system. The network need have no prior knowledge of which users may wish to connect to which other users or resources in the net. This is particularly important (and necessary) for the mobile subscribers.

C. Mobility

The system should support mobile terminals and computers at normal vehicular ground speeds within the area of coverage. Packet radios in mobile applications must satisfy reasonable size, weight, and power consumption constraints.

D. Internetting

The packet radio network structure should be capable of internetting in such a way that a user providing a packet address in another net can expect his network to route the associated packet to a point of connection with the other net or to an intermediate (transit) net for forwarding. Similarly, arriving internet packets should also be routed to the local user.

E. Coexistence

Radio frequency characteristics of the packet radio system should allow coexistence with existing users of a chosen frequency band. This could provide a greater degree of spectrum sharing, particularly among similar systems, and may facilitate the introduction of the technology in new geographic locations.

F. Throughput and Low Delay

The capacity of the packet radio system should allow for variable length packet sizes up to a few thousand information bits, and provide delivery of packets with delays on the order

of 0.1 s in nets of 100 mi area coverage size. Parameters on this order are required to provide the real-time interactive services, and to accommodate efficient data transfers. With 100 kbit/s signaling rates, a maximum packet size might be a few thousand bits.

G. Rapid and Convenient Deployment

Deployment of the packet radio net should be rapid and convenient, requiring little more than mounting the equipment at the desired location. No alignment procedure should be required, and in most applications omni-directional antennas would be used, thus eliminating the need for antenna alignment. Once installed, the system should be self-initializing and self-organizing. That is, the network should discover the radio connectivity between nodes and organize routing strategies on the basis of this connectivity and on the source/destination data of traffic presented to the net. Packet radio should be capable of unattended operation.

H. Error Control

Data integrity is crucial for most computer applications. Error control should be provided by the network, so that packets delivered to a user with undetected errors occur less frequently than about one in 10^{10} packets. This is a critical requirement for computer communications, since even one undetected error in a large file may render it useless or cause troublesome and unpredictable problems during subsequent use of such a file. While detection of errors is essential, choices exist in dealing with the detected errors. In some cases, error detection and retransmission may be used, while in other environments, more sophisticated forward error correction technology must be used in order to maintain satisfactory throughput and delay when communicating through land mobile radio channels.

I. Routing Options

The network should support efficient communication between any pair of users and the capability for users to broadcast a packet to a subset or to all users on the net. Land mobile radio traditionally has been used for point-to-point voice communications. Dispatching services, walkie talkies and, recently, CB radio all have supported a broadcast mode of operation as well. These capabilities could be requested by a "type of service" field provided by the user in the packet header.

J. Addressing Options

The network should provide a capability for addressing a subset of the network participants and for efficiently establishing communication among them. This might be used for real-time conferencing or to support message delivery to a list of addressees with the minimum number of network transmissions. Logical network connectivity is a necessary foundation upon which protocols for these services can be built.

K. Tactical Applications

In tactical military applications, the RF waveform used by packet radios should provide resistance to jamming, spoofing, detection, and direction finding. In many cases, waveforms with these capabilities lead naturally to the capability for position location and relative navigation. With the addition of

communications could then provide identification

III. SIC

Packet radio (air, seaborn, ground) focus on ground, most difficult environment. Terminals are in received signal: man-made structures to multiple the different. As a result of the predict and mobile terminal: packet radio system management capability as a function of multipath phenomena, the use of spread spectrum. In the (roughly) the range for ground characteristic and man-made noise. Discussion of spillover to ground.

Frequency B: The operation is a major in the highest frequency range are determined and propagation requirements.

The required lowest desirable and effective range of RF bandwidth about 0.3. Center frequency excess of the bandwidth would be due to sky. Due to delivery on the system must be second, which is about 300 Hz. From center frequency to the lower high.

Propagation atmosphere, but with spreading factor, multipath spreading factor of 30 MHz.

parameters on the interactive transfers. With these might be a communication security function, the packet radio net would then provide an integrated communication, navigation, and identification system for secure tactical use.

III. SIGNALING IN THE GROUND RADIO ENVIRONMENT

Packet radio technology is applicable to ground-based, airborne, seaborne, and space environments. In this paper, we focus on ground-based networks which encounter perhaps the most difficult environment in terms of propagation and RF connectivity. Ground radio links, particularly when mobile terminals are involved, are subject to severe variations in received signal strength due to local variations in terrain, man-made structures and foliage. In addition, reflections give rise to multiple signal paths leading to distortion and fading as the differently delayed signals interfere at a receiver [13]. As a result of these phenomena, RF connectivity is difficult to predict and may abruptly change in unexpected ways as mobile terminals move about. An important attribute of a packet radio system is its self-organizing, automated network management capability which dynamically discovers RF connectivity as a function of time for use in packet routing. The multipath phenomena also provide a strong motivating factor for the use of spread spectrum waveforms in packet radio systems. In the paragraphs which follow, we first bound (roughly) the radio frequency choices which are most appropriate for ground-based radio networks. We then discuss characteristics of propagation path loss, multipath effects, and man-made noise at these frequencies, and conclude with a discussion of spread spectrum signaling [14] and its applicability to ground-based packet radio systems.

Frequency Band

The operational characteristics of the radio frequency band have a major impact on the packet radio design. The lowest and highest frequencies which can be used for a packet radio system are determined primarily by considerations of bandwidth and propagation path loss (and the associated RF power generation requirement) respectively.

The required systems bandwidth effectively determines the lowest desirable radio frequency in two ways. Practical, off-the-shelf radio equipment is difficult to achieve if the ratio of RF bandwidth to RF center frequency is much larger than about 0.3. This lower bounds the range of acceptable RF center frequencies. In practice, a center frequency well in excess of this lower bound is also desirable if the received signals would otherwise have too wide a multipath spread (e.g., due to sky wave phenomena at HF). For a packet radio system to deliver 2000 bit packets through a network with delays on the order of a tenth of a second, the data rate of the system must be in the range of a few hundred kilobits per second, which implies RF bandwidths of a few hundred kilohertz. From an implementation point of view, then, the RF center frequency should be at least a few megahertz, or in the lower high-frequency (HF) band extending from 3 MHz to 30 MHz. Propagation in the HF band can provide long distance communication due to sky wave reflections from the earth's ionosphere, but the propagation suffers from noticeable multipath spreading of the signal which, as will be described later in this section, limits the data-rate of signals which can be used. Multipath spreading in the very-high-frequency (VHF) band (from 30 MHz to 300 MHz, where line-of-sight propagation

dominates, is typically reduced to a few microseconds as compared to the millisecond spreads encountered at HF, and data rates on the order of a hundred kilobits can be supported. Multipath fading and distortion are still a problem at VHF, particularly for terminals which are mobile or do not operate with radio line-of-sight. However, diversity techniques or the spread spectrum signaling techniques discussed later can overcome these difficulties.

The upper limits on usable radio frequencies for packet radio are primarily established by propagation path loss. As the operating frequency rises to about 10 GHz, absorptive losses due to the atmosphere and rain rapidly increase, and the resulting radio range is reduced accordingly. In general, packet radio systems must use closely spaced relays in order to provide adequate area coverage at these frequencies. The cost of providing a dense relay population may be acceptable if the distribution of users is also dense and if packet radios co-located with the users can provide the relay function. For most applications, however, 10 GHz is a practical upper limit for a useful radio frequency in a ground-based packet radio system. We conclude, then, that practical packet radio systems should use radio frequencies in the upper VHF band, in the ultra-high-frequency (UHF) band from 300 MHz to 3 GHz, and in the lower portion of the super high frequency (SHF) band from 3 GHz to 30 GHz.

An additional factor which must be considered for operational systems is the authorization to radiate packet radio transmissions. The VHF and UHF bands are already heavily allocated. The use of spread spectrum signals potentially could allow coexistence of a packet radio system with existing users of some frequency band. However, this is a relatively new concept from the regulatory point of view, and significant technical issues would have to be resolved to establish the feasibility of coexistence.

The discussion of propagation, multipath, and background noise which follows focuses on the UHF band, although, qualitatively, the phenomena discussed apply to the VHF and SHF bands as well. Later sections of the paper describe an experimental packet radio which operates at 1710-1850 MHz in the upper UHF band.

B. Propagation Characteristics

Packet radio network operations would be greatly simplified if all radios were sited such that a radio line-of-sight path existed to nearby radios. Link design procedures for such paths are well understood, and RF connectivity within the network would then be fixed and reliable. Such stringent restrictions on siting are not reasonable from the user point of view, however. Many users of a packet radio network will have to operate from facilities previously established without consideration of radio propagation. Use of packet radio in a mobile environment would be almost useless if siting were required for reliable operation.

The minimum theoretical path loss is achieved on a radio link in free space (i.e., a vacuum), where received signal strength decreases as the inverse square of link range. For a ground radio link, the path loss of free space may be approached on a link having a radio line-of-sight path, although even under this desirable condition diffraction and multipath phenomena can greatly reduce received signal strength.

When a radio line-of-sight path does not exist on a given link, one can still speak of sited or non-sited terminals, due to

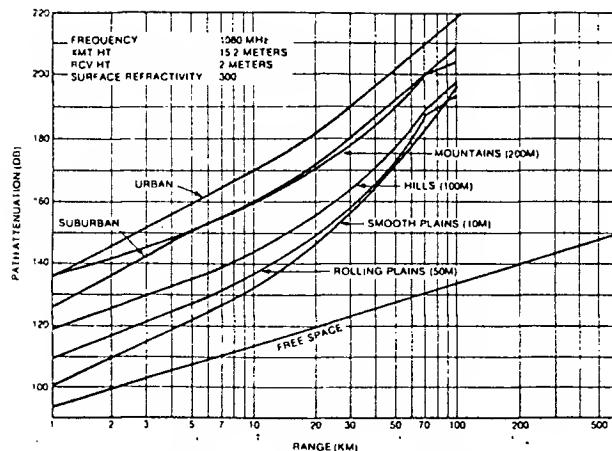


Fig. 1. Path loss versus range. 15-m transmitter height.

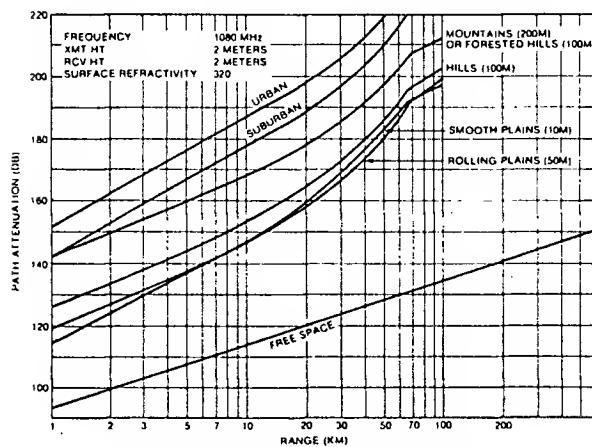


Fig. 2. Path loss versus range. 2-m transmitter height.

the strong influence of shadowing by local terrain and objects and of the elevation of the antenna above the ground. A sited terminal is one which has been located to avoid surrounding obstacles and whose antenna has been elevated to the maximum extent possible, while a terminal operating from a moving vehicle would generally be nonsited.

Average path attenuation exceeds that of a free space radio link by a significant amount in the ground radio environment, depending on the type of terrain and the elevation of the radio antenna. The curves in Fig. 1 and 2 show path loss as a function of link range for a frequency near 1 GHz, and illustrate these dependencies for two different transmitter heights. These curves are typical of propagation of UHF, and the variation of mean path loss as a function of frequency is typically much less than the variations due to terrain at a particular frequency. For example, the mean path loss from 700 MHz to 2000 MHz varies about 8 dB, while path loss at a 20-km range is seen from the figure to exceed that of free space by 25 to 80 dB depending on terrain and antenna heights. Furthermore, the path loss in urban and suburban areas, where many area coverage packet radio net applications might occur, is more severe than that of most natural terrain. The curves shown reflect average values of path loss which apply to a link of a given length which is randomly selected without regard to user siting. Well sited radios will typically encounter less path loss than shown in the curves, while poorly sited

radios will encounter larger path losses [35], [36]. These factors lead to large variations in achievable radio range among users and make RF connectivity difficult to predict in a large mobile user community. The objective of packet radio net design is to overcome this difficulty without placing undue restrictions on allowable user locations. In general, this requires automated network management procedures capable of sensing the existing RF connectivity in real-time and instantly exploiting this connectivity for network control and packet routing.

In addition to the wide variations in path loss, the ground-mobile, nonsited radio channel is subject to the effects of multipath propagation. When several differently delayed versions of the radio signal arrive at a receiver, constructive and destructive interference results. For stationary users, the effect of this phenomenon is that additional attenuation of the signal may be observed when the receiver is located at a point on the ground where the signal interference is destructive. Nulls on the order of tens of decibels may be observed. When communicating users are in motion, or when a multipath component arises from a moving reflector, received signal strength fading is observed as a function of time. The rate of fading is proportional to the velocity of user motion. Movement by a mobile user of only a few meters can cause received signal strength to drop below the threshold of the receiver, thus effectively disabling the link. Radio connectivity to a mobile terminal may change dramatically even for small displacements, and for continuous motion, signal strength may fluctuate above and below the receiver threshold several times during the reception of a packet, causing several short bursts of errors in the data or even loss of synchronization altogether.

We assume the modulation technique used by the packet radio results in a transmitted signal which is structured as a sequence of identifiable segments called symbols, where each symbol is one of a finite set of waveforms. In a simple binary modulation technique, one of two symbols is selected for each transmitted bit, depending on the value of the bit. In a spread spectrum system, the set of symbols may change with time, but a binary system would still choose each transmitted symbol from the set of two, which is in use at that particular time. Typically, a receiver processes the arriving waveform symbols one at a time, making a decision on each one as to which of the finite set has been received. The existence of multipath signal components affects the reliability with which symbol decisions can be made by causing symbol distortion and intersymbol interference. Intersymbol interference occurs when a symbol is overlapped by the delayed components of adjacent symbols. Such interference can lead to lower limit on symbol error probability which cannot be improved by increasing the signal to additive noise ratio on the radio link. When simple modulation techniques, such as phase-shift keying (PSK), are used, the symbol rate must be low enough that only a small portion of the symbol is overlapped by multipath signals of adjacent components. More sophisticated receivers (e.g., using adaptive equalization) can improve performance by suppressing the multipath components, but in order to do this they must rapidly obtain good estimates of the channel pulse response, which can be very difficult if not impossible to achieve. Spread spectrum signals are capable of reducing intersymbol interference effects, as described in Section III-D.

One way to break through the intersymbol interference barrier and achieve megabit rates with binary signaling is to use as

6]. These range among t in a large, t radio net cing undine general, this are capable al-time and work control

the ground- a effects of tly delayed constructive y users, the enuation of located at a e is destruc be observed. a multipath received signal

The rate of sition. More nse received the receiver, ectivity to s for small d strength may several times short bursts n altogether. y the packet ructured as a s, where each simple binary

selected for the bit. In a change with h transmitted that partic ing waveform ach one is to e the existence of ty with which bol distortion ference occurs on components of o lower limits improved by the radio link hase-shift key- w enough that 1 by multipath cated receiver performance is order to do ie channel is not imposible of reduced ed in Section

ference bandw is to use as

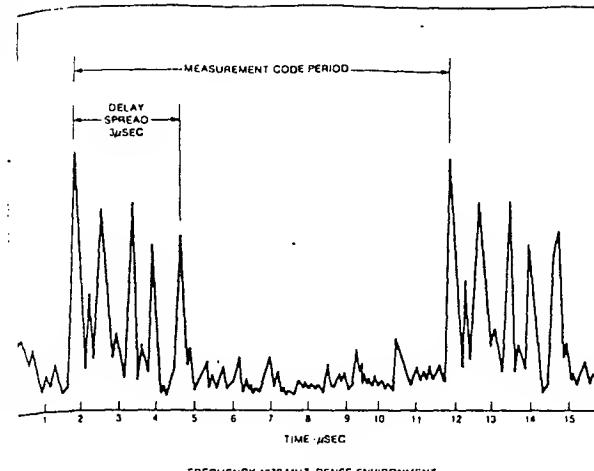


Fig. 3. Received multipath energy.

simulator-correlator type of receiver, such as typified by the RAKE structure [15] and its successors [16], [17]. This type of receiver, which is relatively complex, first estimates the channel impulse response and then adjusts its matched filters accordingly to match the channel characteristics. An alternative approach is to use M -ary, rather than binary, signaling to increase the bit rate to the desired value while still keeping the symbol rate sufficiently low to avoid the intersymbol interference problem. For example, a maximum rate of 250K symbols per second is reasonable to expect with a 4-μs multipath spread; thus four bits per symbol will be needed to achieve a 1-Mbit transmission rate. In this example, the symbol alphabet must contain at least 16 signals, each of 4-μs duration. The M -ary receiver may be relatively complex compared with more conventional schemes such as binary modulation.

In addition to causing intersymbol interference, multipath smears individual symbols by smearing the received signal components over time. Another view of this effect can be seen in the frequency domain. If signal components arrive with differential propagation delay spread dt seconds, nulling of RF signal components occurs at intervals on the order of $\Delta f = 1/dt$. If the spectral width of the signal exceeds the value Δf_c called the channel coherence bandwidth, notches occur in the received signal spectrum spaced B_c Hz apart. A matched filter receiver designed to match the undistorted (or unnotched) symbol spectrum will thus suffer a loss in detection performance since it is no longer matched to the actual received signal. Either the symbol rate must be restricted so that the symbol spectrum falls mainly within the channel coherence bandwidth or other techniques such as spread spectrum must be used to suppress the distortion.

Fig. 3 illustrates the received multipath energy as a function of delay for an urban propagation path at 1370 Hz as measured with a spread spectrum test waveform. Note that significant delay components span a 3-μs interval. The figure illustrates the presence of several major multipath components for each symbol. Ideally, without multipath, the received energy would be concentrated at the beginning of each symbol as a single pulse of short duration. Thus a simple binary modulation system such as phase shift keying would begin to suffer significant intersymbol interference and symbol distortion at data rates exceeding a few hundred kilobits per second.

C. Man-Made Interference

Man-made interference in the RF frequency band includes both intentional and unintentional interference. Resistance to intentional interference (jamming) is of utmost importance in tactical military applications and strongly affects the details of waveform design and system complexity. Unintentional interference results from sources such as automobile ignition spark discharges, arc-welders, electric trains, radars, and ac power distribution systems. These sources of interference can be characterized as impulsive in nature. Such interference is often generated at relatively low signal strength levels, but in some environments, such as urban areas, the density of interference sources is high enough that packet radios will be near enough to some of them to encounter impulsive interference levels which are 60 to 80 dB above the thermal noise level of the receiver. Measurements near major roads and freeways [13] indicates that packets of 10-ms duration would typically suffer interference from several impulses of this magnitude during reception in such areas. Packet radios in such an environment might therefore experience a few bit errors in almost every packet received, and would require forward error correction in order to maintain system throughput.

D. Spread Spectrum Waveform Design Considerations

The packet radio signaling waveform must be designed to perform well with respect to both the natural environment and the induced environment arising from both intentional and unintentional interference including system self-interference arising from the multiple access/random access nature of the packet radio system. We have already noted that the limitations on signaling rate due to multipath can be reduced by using spread spectrum techniques. In addition, spread spectrum provides rejection of interference and the ability to coexist with other signals in the RF band. For these reasons, we consider spread spectrum signaling and identify the performance attributes of two major types of spread spectrum signals which make them well suited for packet radio applications. We do not consider non-spread signaling waveforms in this paper, nor other bandwidth expansion techniques such as FM. It should be noted that spread spectrum is, in reality, nothing more than a particular form of low rate coded transmission and that higher data rates are possible to achieve in principle through the use of more sophisticated systems which couple coding techniques directly with bandwidth expansion. However, engineering and implementation of these coded systems can be quite difficult, particularly at bandwidths in excess of a hundred megahertz. The following discussion illustrates that a variety of spread spectrum systems are well suited to deal with the RF environment, both natural and man-made.

The most commonly used forms of spread spectrum waveforms are direct sequence pseudo-noise (PN) modulation, frequency hopped (FH) modulation, and hybrid combinations of the two. A typical form of PN modulation is illustrated in Fig. 4(a). A source produces binary data at a rate of R bits per second. A pseudo-random generator produces a stream of binary "chips" at nR chips per second where n is an integer which is one or more orders of magnitude greater than unity. The term "chip" is used to distinguish the PN stream, or chip pattern, from the data stream. Each data bit is modulo two added to a sequence of n chips to form a "PN modulated" data stream which is then input to a modulator in order to

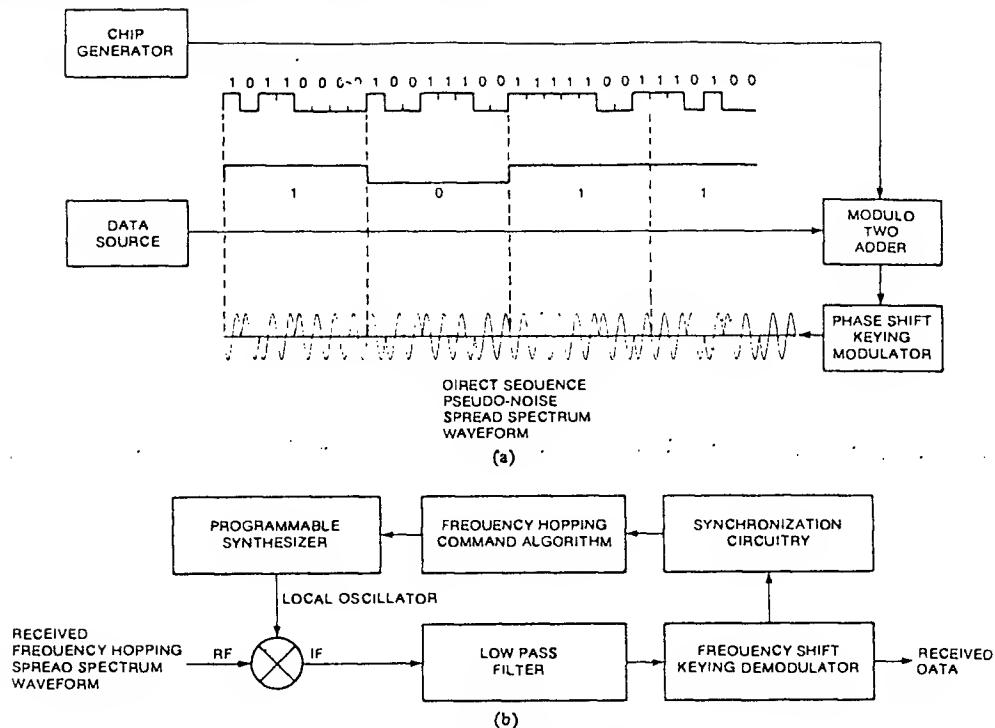


Fig. 4. (a) Pseudonoise spread spectrum modulation. (b) Pseudonoise spread spectrum demodulation.

convert it to an analog form suitable for transmission over the radio channel. The phase shift keying modulator illustrated in the figure is one way of accomplishing this digital-to-analog conversion. A typical form of demodulation is shown in Fig. 4(b). The effect of PN modulating the data stream is to increase the digital rate going into the PSK modulator from R bits per second to nR chips per second. Consequently, the occupied RF bandwidth of the resulting waveform is increased by a factor of n .

Several benefits are received in return for this increased bandwidth.

1) By using matched filter or correlation techniques, the signal-to-noise ratio is improved by a factor called the processing gain which is equal to $10 \log(n)$ dB in the above system. This improvement is realized with respect to both interference and receiver thermal noise as long as the interference in the RF channel is not highly correlated with the chip sequence used to encode a particular bit. In tactical systems, it is necessary to continuously change the chip pattern with time, so that a jammer cannot mimic or replay the waveform to appear as a highly correlated signal at the receiver.

2) The wide-band PN waveform can provide the ability to separate the various multipath signal components using correlation or matched filter techniques. Once separated, these signal components can be combined to reduce signal fading over time and improve signal-to-noise ratio. In addition, a suitably wideband signal will experience frequency selective fading only over small portions of the band. The received energy is therefore relatively constant over time and the overall communication reliability is quite similar to that of a frequency diversity system.

3) Because the signal is spread over a wider bandwidth than would be required for transmission of the signal without PN modulation, the spectral density of the signal is reduced for a

constant transmitted power level. This factor coupled with the pseudo-random nature of the waveform gives the system a lower electromagnetic profile.

4) If the PN chip sequence used changes with each bit, the waveform will have a strong capture property, as explained in Section VI. Capture is the ability of a receiver to correctly receive one packet in the presence of other interfering packets. The capture property greatly enhances multiple access efficiency.

5) By using different PN chip patterns, various groups of users can coexist in the same area with greatly reduced interference. This "code division multiple access" (CDMA) performance again depends on the ability to receive one chip pattern while rejecting others as noise by using matched filter or correlation techniques.

6) Intersymbol interference may be suppressed even if the same chip pattern is reused for each symbol provided that the multipath spread is only a fraction of a symbol duration. If the multipath spread exceeds a symbol duration, intersymbol interference can still be suppressed if chip patterns are changed from symbol to symbol.

Fig. 5 illustrates one of many variations of a frequency-hopped system. Data from the source is input directly to a modulator, in this example a frequency-shift keyed (FSK) modulator is assumed, to produce a signal which might ordinarily be converted to a fixed RF frequency for transmission. In this case, however, a programmable synthesizer is used to select a local oscillator (LO) to dynamically select the RF frequency for transmission. The LO pseudorandomly hops among n frequencies over time as determined by a suitable algorithm. As with PN modulation, the bandwidth used by the system is n times that used without frequency hopping and the overall bandwidth of a FH system yields certain advantages in terms of expanded bandwidth.

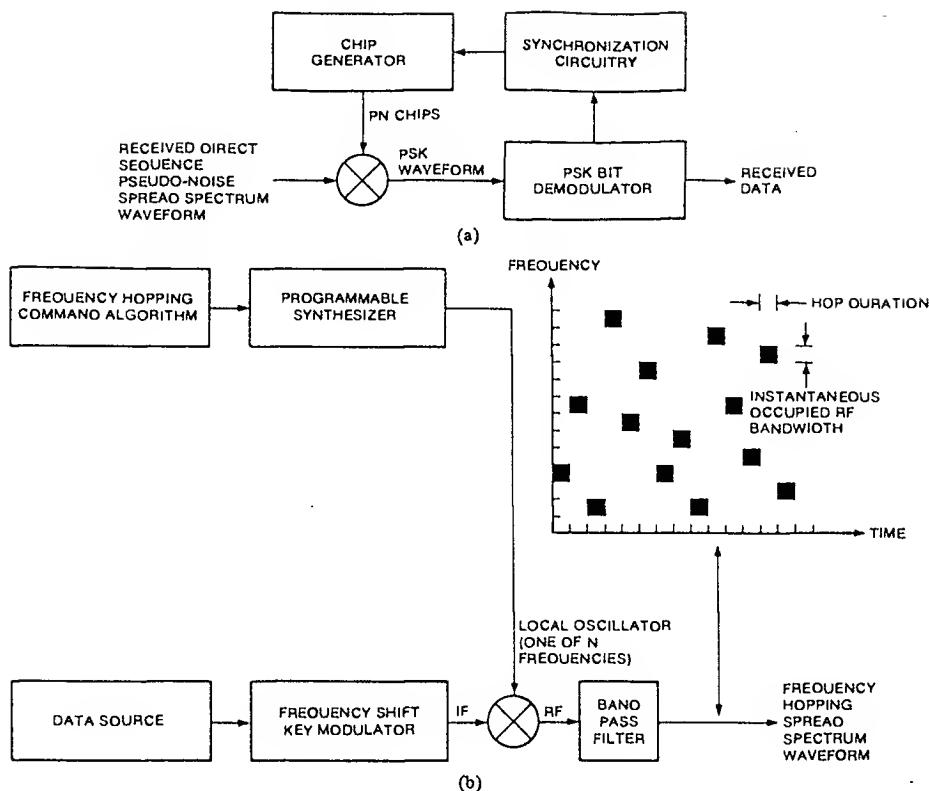


Fig. 5. (a) Frequency hopped spread spectrum waveform generation. (b) Frequency hopped spread spectrum reception.

or coupled with
ves the system

th each bit, the
, as explained in
ver to correctly
erferring packets
multiple access

arious groups of
y reduced inter-
" (CDMA) cap-
one chip pattern
atched filter or

cessed even if the
provided that the
bol duration. In
ion, intersymbol
erns are changed

of a frequency

out directly to a
ift keyed (FSK)
which might not
for transmission
esizer is used as
the RF frequency

y hops among s
uitable algorithm
l by the system
ng and the wide
vantages in retur

1) The degree of interference rejection afforded by a FH system depends on the rate with which the RF carrier frequency changes and on auxiliary techniques such as forward error correction. The basic concept is that data may be received in error when transmission occurs on frequencies subject to strong interference, but that data correctly received on other frequencies subject to weaker interference will allow error correction to restore the lost data. In tactical systems, it is often required that the hopping rate of the system be sufficiently fast that a jammer cannot detect the frequency being used and jam this frequency while it is still being used

because of the time delay accumulated by the signals during propagation over normal operational distances. If the interference is distributed uniformly over all n hopping frequencies, then, as in the PN system, the system may realize an average signal-to-noise ratio improvement of $10 \log(n)$ dB by filtering the instantaneously correct RF band from the total RF bandwidth.

2) FH systems do not inherently suppress multipath fading effects. However, if the average frequency spacing between successive hops exceeds the channel coherence bandwidth, diversity reception may be implemented by sending data once at each of two or more successive hopping frequencies.

3) The coexistence properties of FH are different from those of a PN system. A PN system has reduced spectral density. A FH system also has reduced spectral density, when the spectrum is averaged over many hops. However, the short-term spectral density in any portion of the band will be relatively high during the $1/n$ fraction of the time that this portion of the band is selected for transmission. As a result,

FH systems may interfere with other systems which are sensitive to low duty cycle interference due, for example, to loss of synchronization when the FH signal falls in-band.

4) The capture properties of an FH system are very effective, since the filtering process in an FH receiver may be able to handle interference over a wider dynamic range than is possible in a PN system.

5) If they use different hopping patterns, FH techniques can support overlayed CDMA networks in a common area thus yielding increased throughput in the FH band.

6) If each RF hop is used to transmit exactly one data symbol, FH systems can suppress intersymbol interference. Symbol distortion, however, may still occur if the signaling rate is high enough that the short-term signal spectrum significantly exceeds the channel coherence bandwidth. It is usually necessary to use a noncoherent signaling technique in FH ground radio systems, since the medium may destroy signal coherence from symbol to symbol due to frequency hops that exceed the channel coherence bandwidth.

Combinations of FH and PN techniques can result in a waveform with desirable attributes of both techniques. For example, on each RF hop, one could transmit one of M PN sequences in order to send $\log_2 M$ bits of information. Such a technique could achieve diversity reception, based on suppression of multipath components and reduced symbol distortion, but a receiver for this type of system can be complex.

IV. BASIC SYSTEM CONSIDERATIONS

In this section we introduce the basic packet radio system concepts along with a description of the key elements and how

they work. We describe two key routing options (point-to-point and broadcast) and the protocols which support them. Packet radio combines the use of time division multiple access with radio broadcasting to form a powerful, self-organizing store-and-forward network. We begin with a discussion of the multiple-access channel and its control.

A. The Multiple Access Channel

A multiple access (MA) channel is one which two or more users may nominally share at the same time. A simple form of MA channel is obtained by partitioning the channel into separate nonoverlapping frequency subbands with each user assigned to a separate subband. Another form of MA channel is obtained by scheduling each user's transmission to occur in short nonoverlapping intervals in time. In the first case, known as Frequency Division Multiple Access (FDMA), each user has access to a dedicated portion of the channel at all times. In the second case, known as Time Division Multiple Access (TDMA), each user has access to the whole channel for only a portion of the time. In these two cases, the signaling waveforms are orthogonal in frequency and time, respectively. With a properly designed receiver, orthogonal signals should not interfere with each other in reception. Other desirable forms of multiple access are possible, such as the use of spread spectrum waveforms. Spread spectrum signals may overlap in both time and frequency but use signaling structures that allow a receiver to separate one signal from the others using correlation or matched filtering techniques.

Although relatively simple to implement, a disadvantage of FDMA is that the use of separate dedicated frequency bands between each pair of users does not generally result in efficient use of the frequency spectrum or provide cost effective interconnection strategies for interconnecting multiple users. Only in the event that each pair of users can make almost full use of their frequency subband will this strategy be at all efficient.

We assume each user has a radio receiver that is capable of receiving only one frequency subband at a time and refer to him as a single channel user. The simultaneous use of many radios or receivers per user for increased connectivity is simply not cost effective and the protocols required to support their use would be inelegant and unworkably difficult in practice. If single channel user *A* is communicating with single channel user *B* at a given time, then user *C* cannot talk to user *A* or *B* unless they all share the same channel. Inaccessibility of this form is clearly undesirable in the network environment.

Due to the burst nature of its transmissions, a TDMA system requires time synchronization of its transmissions to achieve nonoverlapping bursts. For this reason, a TDMA system is more complex to implement than FDMA, and some form of local or global control is typically needed to schedule the transmissions so they are nonoverlapping in time. However, an important advantage is the connectivity which results from the fact that all receivers listen on the same channel while all senders in a TDMA system transmit on the same common channel at different times.

A common radio channel provides a cost-effective way to achieve complete network connectivity between subscribers and efficient spectrum utilization can be achieved through statistical multiplexing of many traffic sources onto the channel using TDMA. Each user is guaranteed to be able to communicate with every other user when he transmits at his

scheduled transmission times. In principal, a very large number of users each with very low duty cycle can simultaneously and efficiently share the channel. In order to achieve efficient performance in practice though, a control mechanism is required for access to the channel.

B. Controlling the MA Channel

A variety of theoretical and experimental studies have been carried out to determine the most effective techniques for sharing a MA channel. One of the simplest techniques, known as pure ALOHA was designed for very low duty cycle applications and involves no control other than the ability to detect overlapping packet transmissions (conflicts) when they occur and to randomly reschedule these unsuccessful transmissions at a later time. This scheme is normally implemented using a positive acknowledgment and time-out procedure based on packet checksums. Packets are simply transmitted randomly in time according to the underlying arrival process [18]. A time slotted version of random access is also possible. In this case, accurate timing must be maintained on the channel such that the radios can synchronize their transmissions to begin only at the time slot boundaries. A packet transmission occurs within a time slot with some underlying probability derived from the arrival process. This access method is often referred to as slotted ALOHA [19].

A pure ALOHA system is more efficient than a system consisting of many subchannels dedicated to a collection of very low duty cycle users, but the expected number of conflicts increases as the pure ALOHA channel loading increases until the channel reaches its maximum throughput at $1/2 e$ (about 37 percent) of peak channel capacity (peak instantaneous transmission rate) for fixed length packets and Poisson arrivals. In this same assumed type of traffic, slotted ALOHA achieves twice the maximum throughput of random access (or pure ALOHA). A slotted system may appear advantageous for reasons of efficiency, but the system is also more vulnerable in this case since packet synchronization preambles which are critical for packet acquisition begin in predictable locations in time, and this information could be exploited by an adversary in tactical or military applications. Both pure ALOHA and slotted ALOHA can become unstable and some form of stability control is desirable in operation.

One of the more efficient control techniques for the ground radio MA channel is Carrier Sense Multiple Access (CSMA). In CSMA, each sender first senses the channel, and then transmits a packet only if the channel is idle. If the channel is determined to be in use, the transmission is rescheduled at a later time when the same procedure will be invoked. Various elaborations on the CSMA scheme offer the possibility of achieving 80-90 percent utilization of the channel with low end-to-end transmission delay per packet [20] [23].

In a slotted carrier sense system, each transmitter senses the channel during the beginning of each slot. This elementary observation has further implications when we address the network aspects of spread spectrum in Section VI.

C. System Structure

When critical parts of the system are geographically distributed and often unattended (or unprotected), special provisions are required to ensure the overall network integrity. One way to protect against compromise due to tampering

large number of radios in the field is to control all functions which can have system-wide implications from one or more protected stations. For example, a fully distributed routing algorithm such as that used in the ARPANET would not be suitable for use with unprotected radios, since a small modification at one radio could totally affect the network routing and performance (e.g., imagine that a radio declared itself to be the shortest path to everywhere). Aspects of the network protocols (such as the radio acknowledgment procedures) which must be performed by each radio would be distributed among the radio elements. However, all network control protocols which can have global effect are specifically initiated by one or more entities in the network called *stations*. The resulting network control thus takes the form of a two level hierarchical system. The normal mode of operation utilizes a single station and multiple stations. However, a stationless mode is also possible. The implications of stationless operation are discussed later.

The functions of a station are associated with global management of the radio net [24]. Generally speaking, each station is aware of all operational radios in the network. The stations discover the existence of new radios waiting to enter the network and determine when other radios have departed. The station determines the route to each of these radios and plays an active role in initializing, organizing, and maintaining the operational network. In particular, all routes are assigned by the station to minimize PR cost and complexity. PR's are not required to store information about every other PR and terminal device in the network.

One of the requirements for controlling the PRNET is assessing the reliability of radio links between PR's and using this information to assign good routes. A primary source of this information is the PR neighbor table whose entries are selected by each radio, summarized, and regularly sent to the station along with other status information. For example, each radio reports which other radios it can hear along with raw or processed information for the station to determine the quality of the transmission path between these radios. The station then deduces the overall connectivity of the network (we assume topologically rather than topographically) and determines good routes to itself from each of the radios in this subset. The station then distributes to each radio in its subset the route from that radio to the station. This process is known as *labeling*. The neighbor table is maintained by each radio to determine whether or not an operational station is present and can be used in a stationless mode if necessary.

We assume that a set of radios distributed throughout a geographic area, which we call the backbone, provides a carrier-based packet communication network service to the users. These backbone radios, known as repeaters, receive packets from nearby users and relay them. The repeaters also accept packets from other nearby repeaters for relaying. This extends the range of the system beyond line of sight. Users communicate with each other via a common frequency band [3].

For military operation, where a separate backbone network might be infeasible to deploy, each user's radio might be equipped to support not only his own traffic but that of other designated users. That is, the user's radio may also have "double up" as a repeater, to support network traffic. In this case, we do not identify a separate backbone repeater network *per se*, since it would be indistinguishable from the network of user packet radios.

aphically, special network integrity to tamper

D. Store-and-Forward Operation

An individual packet radio unit is a small piece of electronic equipment which consists of a radio section and a digital section which controls the radio [26]. The radio section contains the antenna, RF transmitter/receiver, and all signal processing and data detection logic associated with modulation and demodulation. The digital section contains a microprocessor controller plus semiconductor memory for packet buffering and software. The radio and digital sections are connected by a high speed interface (see Fig. 6). For each transmitted packet, the digital unit selects the transmit frequency (normally fixed), data rate, power, and time of transmission. In addition, it performs the packet processing to route the packet through the network. In a half duplex mode of operation, a radio may be transmitting or receiving, but not both simultaneously. In the remainder of this paper we assume that each radio operates as a half duplex transceiver in the common frequency band.

Normal store-and-forward operation within the network takes place as follows. A user generated packet with associated addressing and control information in the packet header is input to the digital section of his packet radio, which adds some network routing and control information and passes the packet to the radio section for transmission to a nearby repeater which is identified within the packet. Upon correct receipt of the packet, the nearby repeater processes the header to determine if it should relay the packet, deliver it to an attached device, or discard it. Several nearby repeaters may actually hear the packet, but only one repeater (which we call the next downstream repeater) will typically be identified to relay it. The other repeaters will discard the packet. The packet will then be relayed from repeater to repeater through the backbone (in a store-and-forward fashion using the procedure described above) until it arrives at the final repeater which broadcasts it directly to the user's packet radio. At each repeater, the packet is stored in memory until a positive acknowledgment is received from the next downstream repeater or a time-out occurs. In the latter case, the packet will be retransmitted. Each packet is uniquely identified by a set of bits in its header called the Unique Packet Identification (UPI). The relaying of a packet by omnidirectional broadcast can also serve as a reverse acknowledgment (an echo acknowledgment) to the previous upstream repeater as long as the system is nonspread or a common spread spectrum code is used by the radios.

The propagation delay between radios in the ground radio environment is typically on the order of a small fraction of a millisecond (it could be as low as a few microseconds). As a result, the minimum time a packet must be stored in each digital unit awaiting a positive acknowledgment when it was correctly received downstream is little more than the time to transmit the packet and for the next downstream repeater to quickly return the acknowledgment. Each digital section therefore needs to contain enough memory to store at most a few packets (e.g., 4-6). Due to the half-duplex operation, a repeater which has just transmitted a packet should, under normal conditions, immediately enter receive mode to listen for the acknowledgment.

If, instead, another packet should arrive during the time-out period at the packet radio waiting for the acknowledgment (due to the half duplex operation, the radio must be in receive mode to receive the acknowledgment), the radio will not accept

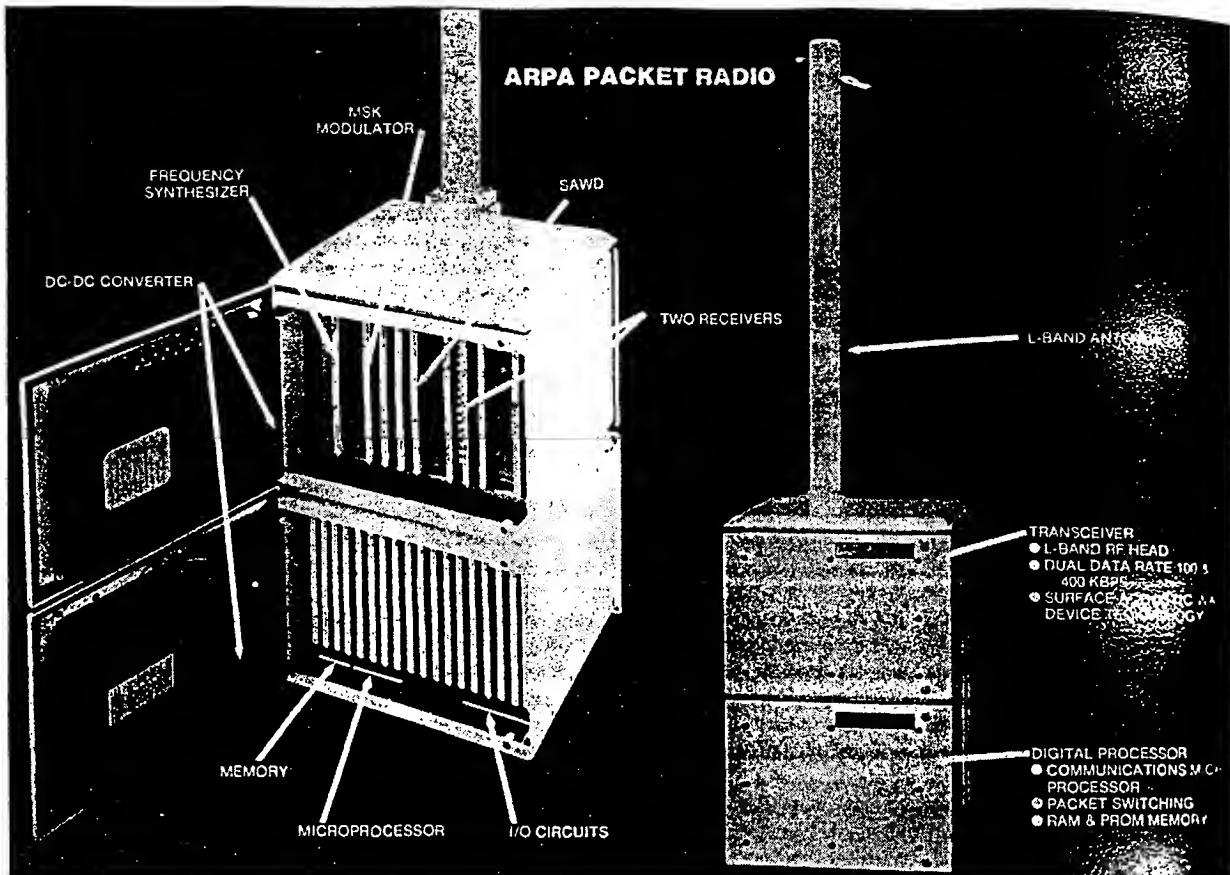


Fig. 6. An experimental packet radio.

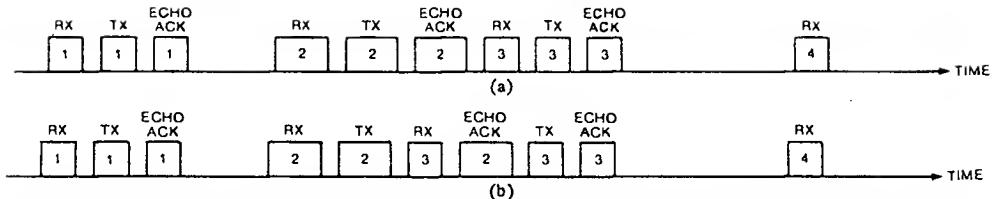


Fig. 7. (a) Normal packet relay cycle (carrier sense mechanism may have delayed Packet 3 until completion of acknowledgment for Packet 2). (b) Alternative packet relay cycle (carrier sense mechanism delays echo acknowledgment for Packet 2 until packet 3 is received).

it if there are no unoccupied buffers in its memory. The half-duplex operation is keyed to a nominal cycle of packet transmission, acknowledgment receipt, new packet receipt, followed then by its transmission and acknowledgment (see Fig. 7). By providing only a few buffers in each packet radio, efficient system performance can be achieved with half-duplex transceivers[27]. A large backlog of packets never accumulates within the individual radios.

Along with mobile operation, rapid deployment and portability are two of the essential attributes of the packet radio technology. Reliable system operation is achieved under these constraints in a number of ways. Each packet radio unit shares a single wideband channel with omnidirectional antennas and extensive siting or alignment procedures are therefore unnecessary in practice. Good connectivity can be maintained with mobile terminals as long as a line of sight path exists. We assume that with omnidirectional antennas (in the azimuthal

plane) communication is possible with a multiplicity of radios within approximate line of sight propagation range. It is further assumed that the network traffic matrix cannot be estimated accurately or is expected to change very rapidly. In this case, dynamic allocation of a common frequency band is a very good strategy.

For the following discussion, we refer to the operation of an experimental packet radio, in which a transmitted packet has the structure shown in Fig. 8. It consists of a 48 bit preamble followed by a variable length header (typically 96-144 bits) followed by the text and a 32 bit checksum.

The packet preamble is used by the radio section of the receiver for several purposes. The first few bits are used to detect the carrier energy and to set the automatic gain control (AGC) to compensate for differing signal strengths of the arriving packets. Correct reception of the packet is totally dependent upon acquisition of the preamble. The next section

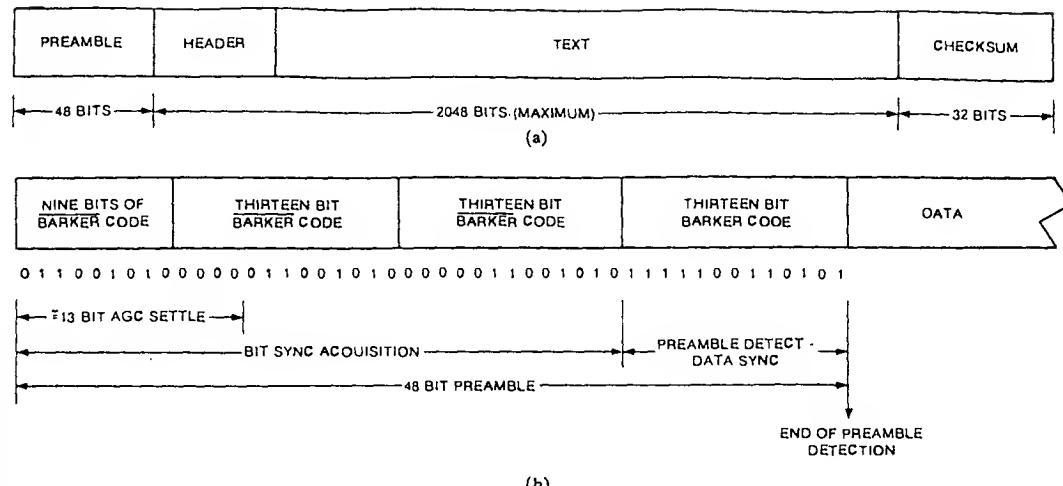


Fig. 8. Structure of a transmitted packet. (a) EPR packet format. (b) EPR packet preamble detail.

are used to acquire bit timing. Following these, the next 48 bits is used to acquire packet timing (identify the end of the preamble and the start of the header). Both the header and text are delivered from the radio section to the digital section which knows the header format and can therefore determine the exact start of the text. The error control bits consist of a checksum appended by the transmitter and checked by each receiver. After checking, the error control bits are stripped off by the radio section as was the preamble before it. The digital section checks a status register in the interface to determine if the packet is correct.

The per packet error rate before error control is usually much higher in the ground radio environment than over a leased line. A packet of duration 5-10 ms will typically experience 2 or more bit errors due to automobile ignition noise and other sources of interference. In addition, mobile terminals experience signal strength fading while in motion due to shadowing and destructive multipath interference which are described in Section III. As a result, and depending on environmental conditions and type of error control in existence at the time, it is possible that a packet will not be successfully relayed even after several attempts. Two actions are appropriate for the system to take in this case. The packet may be dynamically rerouted via an alternate repeater, or an end-to-end retransmission strategy may be invoked external to the network. In the latter case, the packet may either come back at the same repeater at a later time when the original problem is no longer present, or a different route may be provided from source to destination. In the meantime, local network resources (e.g., buffer, channel) are available for more productive use.

Each radio has an identifier which we shall call its selector.¹ The selectors play a central role in the network routing and control procedures. We discuss the point-to-point and broadcast routing options below. The more advanced network control functions are discussed in the next section.

¹For the purposes of this paper, we assume the selectors are unique and preassigned, but this is a simplifying restriction which is adopted for clarity in exposition. Nonunique selectors can, in fact, be used effectively when the probability of nonuniqueness in any given area is sufficiently low.

E. Point-to-Point Routing

In the point-to-point routing procedure, a packet originating at one part of the network proceeds directly through a series of one or more repeaters until it reaches its final destination. The point-to-point route (which consists of an ordered set of selectors) is first determined by a station which is the only element in the net that knows the current overall system connectivity. Having determined a good point-to-point route, where should the station send the point-to-point routing information? One possibility is for it to distribute the information to the individual repeaters along the point-to-point route. In this case, each succeeding packet would only require some form of source and/or destination identifier but would not have to carry the entire route in its header. Alternatively, the station can send it directly to the digital section of the sender's (or receiver's) packet radio. In this case, each packet originating at that radio could then contain the entire set of selectors in its header. However, this choice may have a significant impact on the network efficiency and ultimately its extendability since the selectors would contribute overhead to the packet and, at most, only a small finite set of them could be carried along.

A far more attractive choice is for the sender's or receiver's digital unit to take on the responsibility for setting up the route that was specified by the station. In particular, it transmits a route setup packet (along the designated route) which contains the entire ordered set of selectors and which is used to create appropriate routing entries in the specified repeaters. Regardless of how the routing entries are finally created within the repeaters along the point-to-point route, it may still be desirable to carry along within each data packet the selector for the next downstream repeater, or even the next few repeaters. The latter strategy may have significant operational as well as performance advantages as is discussed further in Section V.

In the event that one or more stations are available in the net and a point-to-point route fails (e.g., an intermediate repeater fails) the existing traffic on that route can be diverted to the station for forwarding to the final destination. The station must recompute a good route to the destination before it can forward the packet. Alternate routing is per-

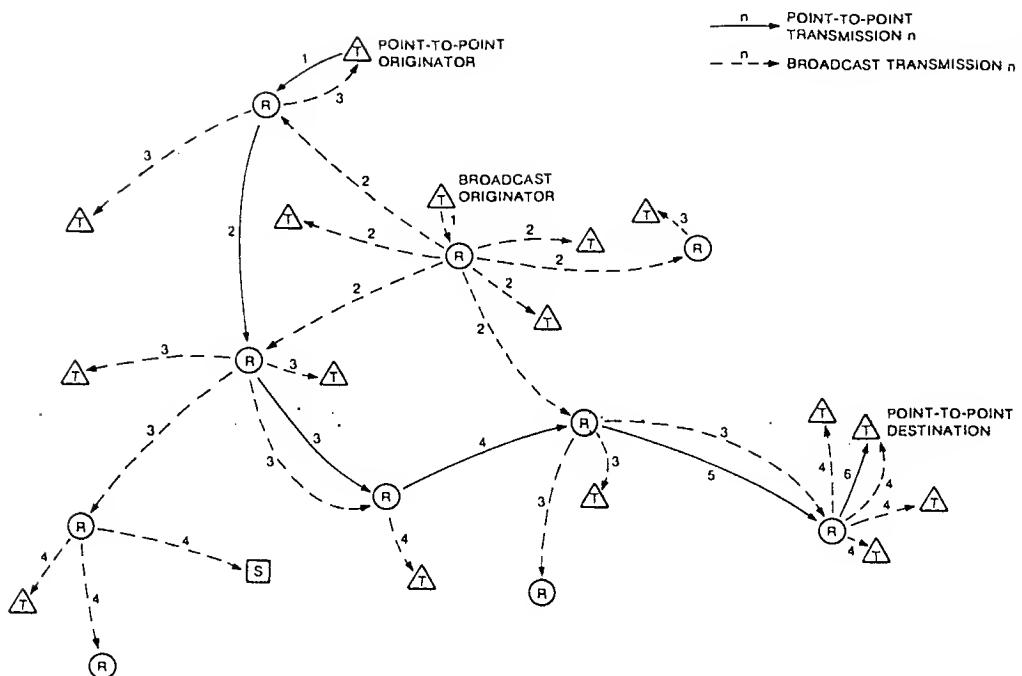


Fig. 9. Point-to-point and broadcast routing.

formed by the PR, when necessary, using parameters assigned by the station. Multiple receivers are allowed to forward a packet around a failed temporarily inaccessible repeater.

As network conditions change (terminal movement, repeater failure or recovery, changes in hop reliability, and changes in network congestion) routes will be dynamically reassigned by the station to satisfy the minimum-delay criteria. In general, only major status changes will result in dynamic routing changes. Hosts and terminals attached to the PRNET remain unaware of (do not participate in) the station's assignment and reassignment of routes; these routes are stored in their attached PR. Establishment of a point-to-point route is an automatic process, transparent to the host or terminal initiating the communication. If the source device's PR does not already have a good point-to-point route to the desired destination, it forwards the incoming packets to some station. The station reroutes the packets to their destination, and in parallel, computes a good point-to-point route from the indicated source PR to the destination. It then deposits this route in the source PR, and later packets are routed directly to their destination instead of being routed through a station.

F. Broadcast Routing

Broadcasting to all radios in the network is another important mode of system operation. In this mode, every repeater in the network keeps a short list of unique packet identifiers (UPI's) for previously broadcast packets that it recently received and retransmitted. A broadcast packet is identified by a bit in its header. If a repeater receives a broadcast packet whose UPI is already on its list, it will discard it. Otherwise it will accept the packet if correct and retransmit it. In the broadcast mode, each packet radiates away from the source radio as in a wave-front type of propagation. The use of a list of UPI's in each repeater prevents any packet from returning to portions of the

net through which it (or any of its copies) has already passed.

Broadcasting is not a particularly efficient mode of operation for two party communications, but it is a very robust way to distribute packets to all parts of the net. This type of service may also be unreliable in the sense that not all radios may actually receive every broadcast packet and end-to-end procedures must be invoked if reliable delivery is essential.

A broadcast routed packet will generally be intended for actual delivery to all users or a subset of users (e.g., conferencing). However, a user could also use this routing option if, for example, there was no working station on the net to supply a point-to-point route to a single intended recipient. In this case, the destination address is used to identify the intended recipient so that it will accept the packet and the others will discard it. Point-to-point routing and broadcasting to all users are illustrated in Fig. 9. The circles indicate repeaters, the triangles represent user terminals and the square boxes denote stations.

G. Mobile Operation

All elements of the packet radio network can be mobile, within a fleet of ships, or certain elements can be fixed while others are in motion. In general, we assume that stations and repeaters move relatively slowly, if at all, so that topological changes in the backbone are not frequent events. Normally, user terminals will also move slowly enough that assigned point-to-point routes will remain effective for at least a few minutes, if not longer. However, certain user terminals may move sufficiently rapidly relative to the other elements in the system that the use of point-to-point routes would not be a practical choice. A particularly significant aspect of broadcasting is that only the identity of the intended recipient need be known to the sender and not their actual location or the routing. Thus broadcasting is a useful technique to

the need for assumed to u high speed fundamental the presenc in real-ti system can ada path, wil mobile terminal significant point routes ir currently servic gically disapp ends as the ance upon s route) provide n. Howeve and the ected by th m after whi point route ll help to sp contact w distress messa to the station point route a mobile unit.

Reliable Delivery
 The inherent unreliability of an end-to-end internal network maps one ou highly efficient net user tra per forma assume an technique to be individual pack work could uti [38], which user's termi device kn forms the er terminal (e. directly to Within the PR control packets each radio continued availa nation cou and remo consecutive lation to are used t as the retr retrans the reliabl basic ele described of syste

less the need for control of rapidly changing routes. A station is assumed to use broadcast routing to communicate with the few "high speed" users.

Fundamental performance limits are imposed on the system by the presence of mobile terminals. The need to update routes in real-time as users move, and the time in which the system can adapt to environmental changes along the mobile user's path, will determine a maximum allowable speed for mobile terminals which use point-to-point routes.

A significant problem associated with the use of point-to-point routes in mobile operation arises from the repeater currently servicing a mobile vehicle's traffic. The radio link can quickly disappear and reappear on a time scale of a few seconds as the vehicle moves past various obstructions. The reliance upon summary reports from each radio (e.g., twice a minute) provides a simple way for station to track the mobile units. However, loss of communication between the mobile unit and the repeater handling its traffic would only be detected by the station after a lapse on the order of a minute, after which the station could assign a different point-to-point route for that mobile unit. The following procedure will help to speed up the process. Whenever a mobile unit loses contact with its next downstream repeater, it broadcasts a distress message, and all repeaters within range will forward to the station. The station can then assign a new point-to-point route and quickly restore communication with the mobile unit.

Reliable Delivery Mechanisms

The inherent undependability of a mobile radio channel requires an end-to-end protocol to provide reliable operation. For internal network control traffic (nonuser traffic), where perhaps one outstanding unacknowledged packet is sufficient, a highly efficient but specially designed protocol suffices. For internet user traffic, or internet operation, a more flexible and other performance protocol is desirable.

We assume an end-to-end error detection and retransmission technique to be used in the network for reliable delivery of individual packets. Each source/destination pair on the network could utilize an end-to-end protocol such as described in [38], which also supports internetworking. In this case the user's terminal could be equipped with a microprocessor-based device known as a Terminal Interface Unit (TIU) which performs the end-to-end protocol, and any local support for the terminal (e.g., local echoing, formatting). The TIU interfaces directly to the digital unit of a packet radio.

Within the PRNET, stations and radios need to communicate control packets reliably. For example, the regular reports from each radio to the station are used to validate the radio's continued availability. Without an effective recovery procedure a station could declare a perfectly good radio to be out of order and remove it from service if several of its reports were to be consecutive. Similarly, parameter change packets from a station to the radio should be delivered reliably since these are used to set dynamically important radio parameters, such as the retransmission interval or the maximum number of allowed retransmissions. The Station-PR Protocol (SPP) provides the reliable delivery mechanism.

V. ADVANCED SYSTEM FUNCTIONS

The basic elements and operation of a packet radio network were described in the previous section. The more advanced aspects of system operation and control are briefly highlighted

in this portion of the paper. We begin by discussing the way in which the network is initialized and the techniques used to introduce new radios and stations into the net. We then discuss the need for multiple stations in a single radio net and describe the basic multistation operation.

The use of stations has clear efficiency advantages when elements of the network are mobile or otherwise prone to rapidly changing fluctuations in communications reliability or connectivity. However the network should support a stationless mode of operation in which no operational stations are available. A method is discussed for implementing good routes under the stationless mode. Finally, we discuss several alternatives for achieving multi-destination routing.

A. Network Initialization and Control.

Consider a collection of geographically distributed packet radios, each of which is powered on and capable of communicating packets to some subset of radios within line of sight propagation range. We assume for the moment that no operational stations are present. Under these circumstances, each radio periodically (e.g., every 30 s for fixed radios and every 5 s for mobile radios) broadcasts a "radio-on packet" (ROP) announcing its existence and containing selected status and identification information from its digital unit. This information is also sent out over its host interface. Some set of radios within range will hear this ROP, will note the event in its tables along with the measured strength of the received signal, and will discard the ROP. There are several ways a PR can determine the quality of a radio link using ROP's. For the purposes of this paper, we simply assume that each ROP contains a cumulative count of the number of packets the PR has received from every other PR it can hear. Upon receipt of an ROP, a radio can determine its connectivity and the percentage of packets successfully communicated on each link. In this fashion, each radio is able to determine the set of local radios with which it can communicate reliably.

If there was an operational station on the net, the PR's would send summary ROP's directly to the station at appropriate times (on a point-to-point route) to convey the current labeling information and also the neighbor table information. This is done both periodically and upon detection by the PR of a possibly significant change in some link. To coordinate transmit counts with receive counts (which is essential for knowing what percentage of traffic on the link is heard), and to get the information to the station, PR's include their transmit counts in each summary ROP, and each receiving PR includes the matching receive count in the summary ROP before forwarding it to the station. All traffic counts are cumulative modulo a very large number.

If the radio is not aware of the existence of a station, it does not transmit any summary ROP's. In this case, it will support the stationless mode of operation described later in this section.

When a new radio is powered on, it will begin broadcasting ROP's, be recognized by neighboring PR's and quickly enter the system. If a radio moves, its new connectivity will be determined locally by the ROP mechanism. If a radio is powered down, fails, or can no longer communicate with any radios, its new connectivity state can also be determined locally (the lack of connectivity can be inferred by the other radios if they keep a record of the previous state or set of states). The ROP mechanism is designed to insure that each radio directly reports its own status and identification information to other radios within range and to the station.

B. Station Entry

When a station is first connected to an operational packet radio, it will soon hear an ROP from that radio over the host interface after which the station will promptly "label" the radio. The labeling process consists of first determining and then supplying the radio with a route (i.e., a set of selectors) to the station. In this first step, the task is trivial since the route to the station is via the host interface. In principle, once labeled, the packet radio periodically will continue to send its ROP's over the host interface to the station along with the less frequent summary ROP's. From the status information in these ROP's, the station will learn which radios are in direct communication range of its radio.

We adopt the following terminology for ease in explaining the labeling process. The radio directly connected to the station is said to be at level 0 with respect to its station. All radios in direct communication range of the level 0 radio and not directly attached to the station are said to be at level 1 with respect to the station. Similarly, a radio in communication range of a level $n-1$ radio which is not itself already at level $n-1$ or less is said to be at level n . A radio may be at different levels with respect to different stations, and its levels may change during mobile operation. A level simply indicates the minimum number of radio hops to the station and does not otherwise affect packet routing.

Having labeled its level 0 repeater, and learned of the level 1 repeaters, the station then reliably labels the level 1 repeaters one at a time using the SPP protocol over the radio channel. The level 1 repeaters then begin sending status reports to the station which prompt the station to label the level 2 repeaters and so forth. Various implementation steps can be taken to insure that labeling occurs quickly and that the station is not inundated with status reports during the labeling process or afterwards. If the order of events is reversed and a packet radio is attached to an already operational station, the sequence of actions taken is identical.

In the case where a station is connected to a network which has already been labeled by one or more other stations, the sequence of actions taken by the station is also identical. Each radio has table space (labeling slots) for storing routes to several stations and by design the number of slots in each radio limits the number of stations that are allowed to label it. The route from a radio to a station is supplied by that station via the labeling process. Two stations which have labeled a common repeater are known as *neighbors*. Summary ROP's, which are sent by each radio to the stations that have labeled it, include sufficient information about the current labeling slot entries for each station to learn which other stations are its neighbors. No direct station-to-station coordination is required to acquire this status information.

One having labeled a radio, the station must relabel the radio within a given time or the labeling slot entry will expire. These entries are timed out (relatively slowly) by the radio, and the age of each entry is reported in its status reports. A station will always fail to successfully label a repeater which has no available labeling slots, but it can refresh an existing entry of its own at any time. A slot whose entry has expired is not erased by the radio (the route is not normally used either) but it may be overwritten by another station if no other labeling slots in the radio are free. In principle, a radio could be provided with diverse routes to a given station for applications such as mobile hand-off.

This portion of the multi-station design achieves the following goals.

- 1) It provides complete redundancy among multiple functionally identical stations in such a way that "hot switch over" is provided, i.e., other stations will automatically and immediately assume the responsibilities of any station which fails.
- 2) It shields all end devices from any knowledge about stations. In particular, no terminal or host device needs to know which stations are currently operational, where they are or even how many of them are in operation. This comment also applies to the single station design.
- 3) It keeps PR's simple for reasons of reliability and economy.

As we shall see below, stations may collaborate in establishing routes between users which are within the jurisdiction of different stations, for graceful handover of responsibility for mobile terminals, and for supporting expansion of the PRNET to several hundred stations and several thousand PR's. After a brief discussion of renewal points and routing, we discuss the multiple station operation in greater detail.

C. Renewal Points and Routing

A renewal point is a PR along the route of a packet where the route (as specified in its header) may be altered. In point-to-point routing, the header contains fields which identify the next few designated repeaters along the path to a specified destination. Every repeater on a point-to-point route can act as a renewal point where these fields are rewritten. Two reasons for identifying more than just the next repeater in the header are to allow a detour via an alternate repeater in the event of failures along the designated path and to allow some repeaters not to serve as renewal points. The detour must be such that it eventually rejoins the designated path.

To function as a renewal point for a point-to-point route, a repeater must have a renewal table containing the next few designated repeaters for that route. When a packet arrives at the next downstream repeater for relaying, its routing fields are rewritten in the header according to the current renewal table entries. To conserve table space, each repeater maintains at most, one table entry for each source/destination PR pair. In addition, the table must identify the last few upstream repeaters on the path so that the source can be notified in the event of communication failures at any point of the path.

A packet may also contain an entire path of route selectors in the text of the packet. This case may be distinguished by a special bit in the packet header which indicates that the text contains all the route selectors. Such a packet is known as a route setup packet and is used to initialize or refresh the renewal table entries in each repeater. Upon receipt of a route setup packet, a repeater extracts the renewal table information (normally a few entries) from the entire list of selectors in the text and writes it into the renewal table. The contents of a route setup packet are normally inserted by the digital unit of the destination PR. Any packet may be a route setup packet, subject only to the maximum packet length constraints of the network. A route setup packet may also contain data.

D. Multiple Station Operation

A single station in a network with a relatively small number of radios (tens to hundreds) is sufficient to organize it efficiently and to control its operation effectively. If the stations

uses the following:
long multiple
at "hot switch"
automatically and
/ station which

knowledge about
device needs to
where they are

This communi-

cy and economy,
e in establishing
isdiction of de-
sponsibility for
1 of the PRNET
nd PR's. After
, we discuss the

a packet when-
tered. In partic-
ular identify the
n to a specified
at route can ad-
en. Two reasons
ter in the header
in the event of
v some repeaters
ust be such that

point route, us-
ng the next few
packet arrives at
ts routing field
current retrans-
beater main-
ination PR path
few upstream re-
e notified in the
of the path. The
of route selected
stinguished by the
tes that the net
et is known to the
or refresh the
receipt of a re-
table informa-

of selectors in the
the content of the
the digital unit
route setup pack-
th constraints
contain data

ely small num-
organize it
ly. If the state

one of two cases will arise. The backup case is a station-mode in which communication is still possible through the net but with a lowered overall utilization. This case is discussed in Section V-E below. In the normal case of multiple station operation, the system control responsibilities are easily assumed and if one station fails the others will temporarily take over its functions with little or no degradation in system performance. Thus, a major incentive behind multiple station operation is efficient and reliable system performance in the event of station failure.

However, a single operational station is no longer sufficient if the number of radios becomes too large (thousands or more) and its capacity is otherwise exceeded. A set of distributed multiple stations are required in that case to maintain the system performance without degradation. Multiple station operation is necessary to support the setup and maintenance of routes through many geographically distributed repeaters under the jurisdiction of different stations and, in particular, to handle mobile terminals at the two ends of these long routes. Various alternatives for multiple station operation are currently under investigation. One of these alternatives is presented below.

Each station individually competes for control of radio resources by attempting to label radios which report free labeling slots or expired entries. If there are more labeling slots per radio than stations which are attempting to label them, a maximum number of radios will be labeled by each station. In general, a station has a limit on the maximum number of radios it is allowed to label, as well as on the maximum traffic level it can operationally support. Usually there will be many fewer stations than radios and the station locations will be effectively distributed throughout the net (rather than concentrated) to distribute the station related traffic and to provide greater station coverage of radios in the event that each station cannot label all radios in the net.

We do not treat the case where the density of stations is sufficiently high or the number of labeling slots sufficiently small that there are more stations than slots per radio. In this situation, the stations would greedily compete for the labeling slots and some stations would be prevented from acquiring a set of radios in the labeling process. We also assume that stations do not coordinate with each other to adjust the allocation of radios to stations. By the nature of the labeling process, stations can only label contiguous sets of radios. If a particular entry in a repeater at level $n-1$ should somehow be allowed to expire by its station, and that labeling slot is overwritten by another station, then other radios at level n and higher may also be cut off completely from the station's labeling process. In this case, the labeling slots in all other radios will shortly expire as well.

In the remainder of this paper, we shall assume each radio has a sufficiently large number of labeling slots to accommodate the number of stations which attempt to label it. We also assume that the net is sufficiently large so that each station can label every radio. From the summary ROP's, each station knows the identity of its neighboring stations and a point-to-point route to each radio labeled by itself and a neighbor. If a source and destination are both within the coverage of a single station, that station can handle the route computation. Otherwise, each station must also collaborate with its neighbors to find good routes.

It is an open question as to whether stations should communicate directly with other than their neighbors, except via a

broadcast technique which requires no state information to be kept about the existence of all other stations or how to get to them. One can visualize a logical store-and-forward network of stations interconnected by point-to-point radio routes (the station-to-station routes can also be provided by other media) using distributed adaptive routing. Participation in such a logical network may complicate the station's tasks and will increase the steady state traffic load somewhat on both the station and the radio channel. We only consider the case where neighboring stations communicate directly with each other. The use of a logical network of stations may be a viable alternative if the overall system performance without it is deficient in some significant way.

We next describe routing through the multistation environment. Each station is assumed to know which radios it has labeled, but must inquire of other stations to learn the whereabouts of other radios (and their users). This inquiry is assumed to take place upon request by a user via a broadcast to all stations via neighboring stations. The inquiry process also provides the destination station with a set of selectors for a point-to-point route to the destination station from the originating user (if fixed) or station (for mobile users). If both users are highly mobile, a point-to-point route will be established between the two end stations which will individually handle the final distribution. If both users are fixed, the destination station will choose the last few selectors from among the radios it has labeled to obtain a point-to-point route directly between the two end users (with no intervening stations) and will supply it to the destination user. The actual point-to-point route setup is initiated by the destination user in this case, rather than by the destination station as was the case for mobile users. If one user is fixed and one user is highly mobile, the resulting point-to-point route will be between the fixed user and the remote station which will handle local distribution for the mobile user.

The route selection process takes place as follows. A user's packet radio generates a packet for a destination outside the control of his station and his radio routes it to an appropriate local station (e.g., the closest one according to some metric). The station converts this packet into several distinct route finding packets which it sends to each neighboring station via some repeater jointly labeled by both stations. The conversion involves adding to the packet the station ID and a list of selectors from the user to the jointly labeled repeater. When the packet is received by the neighboring station, it checks to see if the specified destination user is under its control. If not, it again converts the packet into several distinct packets by adding its own ID and a list of selectors from the original jointly labeled repeater to another repeater jointly labeled by a station not previously visited by the packet. If the packet arrives at a station which has just previously handled the same request via another route, it will be discarded by that station. In this way, one or more route finding packets will eventually arrive at a station which has labeled the destination user and will contain a composite list of selectors. The destination station then passes a complete list of selectors to the destination user's packet radio which initiates the route setup procedure described earlier in this section. This process is illustrated in Fig. 10. The route finding path is shown as a series of dotted lines and the point-to-point route is shown as a series of solid lines.

Once having provided a point-to-point route between two

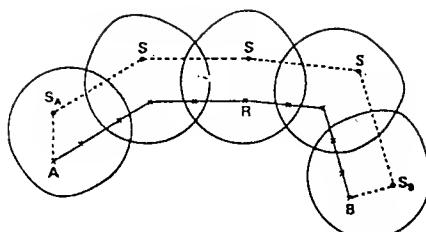


Fig. 10. Route finding and route setup with multiple station operation.

radios under its control, a station does not need to retain detailed state information about that route, regardless of whether the route is part of a larger multistation route or used by two radios under its control. It may be appropriate for the station to retain certain summary information, however, such as the cumulative number of point-to-point routes it has provided, or the number of renewal table entries which were made in each repeater over the last minute, etc.

If a radio fails while a packet is enroute (e.g., repeater R in Fig. 10), and this disrupts a point-to-point route, local alternate routing around the failed radio will take place if possible. The particular station in control of the failed radio will attempt to take a locally corrective measure using only the routing selectors contained in the en route packet (for transit traffic) and the destination ID for nontransit traffic. In this context, transit traffic consists of packets which are passing through the domain of that particular station en route to their final destination. Nontransit traffic consists of packets destined for a radio that has been labeled by that station.

If the locally corrective measure fails (remember there is no state information about specific point-to-point routes in the station), an error message will be returned to the source and the original route setup process will be reinitiated. In this process, certain packets may be discarded, so the end-to-end host protocol must be prepared to recover from this situation.

E. Stationless Operation

The main difference in operation when no stations are present is that each radio initially relies on broadcast routing to all radios in the net in order to communicate with any destination. Broadcasting is a stable and effective way to communicate, particularly with mobile terminals, but it is inherently very inefficient (We assume that distributed adaptive routing cannot be used). For stationless operation, it is highly desirable for the radios to determine acceptable point-to-point routes, even if they only remain acceptable for short periods of time. In a technique that is very similar to multistation route finding, broadcasting can be used to find point-to-point routes which may be usable if the radio links are relatively stable. The repeater and terminals in the network cooperate with each other to discover and set up routes in three phases.

1) *Route Discovery*: A route finding packet is broadcast from the source PR when its attached device attempts to communicate with a destination for which a route is not yet known. This packet contains the source selector, the desired destination selector, and a sequence number which insures uniqueness. Any PR which hears the route finding packet appends its own selector to the data field, stores the information which uniquely identifies the packet, increments a "hop" count in the packet and rebroadcasts it. This PR will then

ignore any subsequent copies it hears of this packet, so the broadcast proceeds outward eventually reaching all PR's in a "wavefront" type of propagation. If the hop count exceeds a maximum value, that packet will be discarded.

If the route finding packet reaches a PR connected to the desired destination, a successful route to the destination is contained in it. Of course, the destination may receive a number of such packets over different routes. With some probability, the route finding packet may also never reach the destination. Radio networks are inherently lossy and no guarantee of delivery can be given. The user can try again or give up. The forwarding of the packet also involves an incremental route delay computation. A delay field is initialized to zero by the source PR; each PR forwarding the packet adds in an estimate of the delay the packet incurred traversing the previous hop. This estimated delay is derived from the PR's neighbor table, which estimates reliability (and hence expected delay, through the hop retransmission algorithm) for the hop from each neighbor.

Once the route finding packet arrives at the destination PR, it contains the estimated round-trip delay for a route which works in both directions. The destination PR will wait long enough to receive most of these packets (which may have travelled over different routes), will select the route with minimum delay, and will store this as the best route.

2) *Route Setup Phase*: For efficient channel utilization, a short routing field in the header is desirable for normal data packets. Since the path discovered above may involve a very large number of selectors, it is clear that not all routing information can be stored in the data packets themselves. Some of the routing information must be stored in the intermediate PR's along the route. This is accomplished by a route setting procedure which is nominally identical to that which is used in the operational station case. A route setup packet is sent from the digital unit of the destination PR which traverses the route specified in the selected route. This packet causes renewal table entries to be written in the intermediate PR's (a table indexed by source and destination). Once this route setup packet arrives back at the source PR, the entire route is set in both directions and we proceed to the final phase.

3) *Normal Data Traffic*: Normal data traffic is sent using the short routing field in the header to specify the next few hops. The source PR and each intermediate PR along the route overwrite this route information with their stored route information for the destination specified in the packet. It is a fresh route for the next few hops each step of the way. As before, the choice of a few next hops rather than a single next hop provides the flexibility of permitting alternate routing if the next PR in the route is failing.

There are four major differences in the functioning of a stationless network and a network with operational stations. First, the station collects connectivity information so that it can supply a route without requiring a broadcast scanning procedure involving all radios; this economizes on channel use and contention. Secondly, the station can compare all possible routes when making its choice, not just the routes which happened to be traversed successfully on a single attempt; this results in a better route selection. Third, the station can detect changes in connectivity quality and issue better routes as appropriate without interrupting on-going communications. Finally, the station can perform global congestion control by changing PR parameters in a coordinated fashion to relieve congestion in different areas of the network. The stationless network [28]



Fig. 11. C

packet, so that all PR's in a count exceeds a limit connected to the destination to receive a PR with some probability never reaches the lossy and can try again involves an interval is initialized during the packet incurred a delay is derived as reliability (and transmission size).

destination PR's, a route which R will wait long which may have a route with intermediate utilization, for normal which involve a very not all routing by themselves. Some the intermediate by a route, which is used in up packet is sent R which traverses his packet can intermediate PR's (since this route sets fire route is set up information to more than one destination (e.g., conference, data base updating, etc.). This capability could be provided by sending multiple copies of the original information, one PR along the heir stored route. The routing capability could be provided for efficiency and the packet delay. We view broadcasting as a special case of multideestination routing where all radios (rather than a subset) other than a receiver participate.

The inherent broadcast mode of the radio channel (the "ether") provides an opportunity for efficient multideestination routing which is not present in dedicated-channel networks. In particular, the "fan-out" required to transmit a message from its source to multiple destinations can be provided by the inherent fan-out of the radio channel. Each transmission can be simultaneously received by many network nodes at no additional cost in radio channel utilization. Alternative strategies for multipoint routing include: Minimal Spanning Tree routing, Source-Based Routing, and Multi-Address Routing. In the first case, one minimal spanning tree is defined for the network, and all traffic flows on that tree to the selected destinations. In the second case, each radio derives its own minimal spanning tree with itself at the root. In the third case, a fashion to of addresses is used to selectively route packets through the network [28].

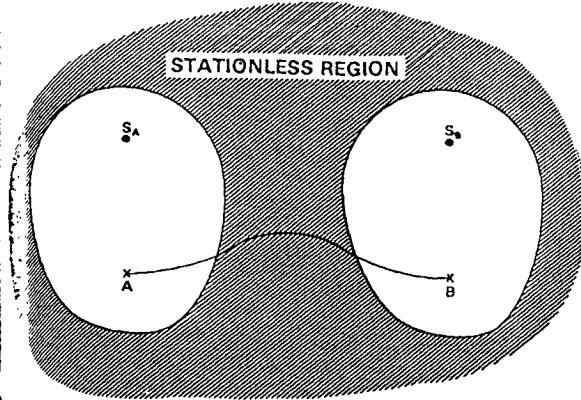


Fig. 11. Communication across a stationless region.

use of PRNET operation is primarily envisioned as a fail-safe (graded) mode of operation rather than the normal mode. There may be instances in which an initial deployment strategy calls for radios to be made operational before stations and enter a stationless mode of operation would be very useful. It is even possible for a set of radios under station control to be within range of another set of radios operating in stationless mode and to communicate with each other. Normally, the stations would attempt to label the other set of radios, if it could accommodate more radios. If, however, it cannot label them for whatever reason, communication can still take place between the two sets by a concatenation of the station routing and the stationless routing. We do not address this area further in this paper, but if two or more station controlled sets of radio are connected via a stationless region, they can communicate by this same procedure (see Fig. 11).

Multideestination Routing

Many applications require the ability to communicate identical information to more than one destination (e.g., conference, data base updating, etc.). This capability could be provided by sending multiple copies of the original information, one copy to each destination. The alternative, a multideestination routing capability could be provided for efficiency and delay. We view broadcasting as a special case of multideestination routing where all radios (rather than a subset) other than a receiver participate.

The inherent broadcast mode of the radio channel (the "ether") provides an opportunity for efficient multideestination routing which is not present in dedicated-channel networks. In particular, the "fan-out" required to transmit a message from its source to multiple destinations can be provided by the inherent fan-out of the radio channel. Each transmission can be simultaneously received by many network nodes at no additional cost in radio channel utilization. Alternative strategies for multipoint routing include: Minimal Spanning Tree routing, Source-Based Routing, and Multi-Address Routing. In the first case, one minimal spanning tree is defined for the network, and all traffic flows on that tree to the selected destinations. In the second case, each radio derives its own minimal spanning tree with itself at the root. In the third case, a fashion to of addresses is used to selectively route packets through the network [28].

A source-based routing scheme is a good choice for broadcast applications and is generally preferable to the use of minimal spanning tree routing. The multiaddress routing scheme is preferred for conference applications involving only a subset of the radios.

VI. NETWORK ASPECTS OF SPREAD SPECTRUM

The use of spread spectrum waveforms in a packet radio system is motivated largely by the desire to achieve good performance in the fading multipath channels resulting from nonsited mobile system users, by the need for coexistence with other systems and for antijamming capabilities in tactical applications. Once a spread spectrum waveform is adopted, however, a number of other benefits, which in themselves suggest the use of spread spectrum, are also available, such as a strong capture capability to enhance access efficiency, the potential for an integrated position location feature, and the ability to operate links, nets, and subnets on pseudo-orthogonal waveforms using CDMA.

These capabilities are not received free of cost, however. Use of a fixed spread spectrum waveform adds a modest amount of system complexity, while a time varying spread spectrum waveform requires some degree of network time synchronization to be established and may require added system protocols to effect the distribution of variables to control waveform generating algorithms.

The use of spread spectrum, although desirable for many applications, is not an *a priori* requirement for a packet radio system. It is worthwhile to distinguish those aspects of network operation which are primarily due to the presence of spread spectrum from those which are not. Consequently, we have identified six spread spectrum network concepts for treatment in this section to separate them from the other more generic packet radio concepts.

A. Synchronization

Because packet radio networks operate with random access transmissions on a common radio channel, a receiver cannot anticipate the sender of any particular packet or its exact time of arrival. The receiver must use the initial portion of the packet (preamble) to detect the arrival of the packet, adjust any automatic gain control loops, and acquire synchronization with the packet symbol and waveform structure to the degree needed to receive the remainder of the packet successfully. In systems which use a spread spectrum waveform which varies with time, packet radios must acquire and maintain a common network time of day, so that the proper reference waveforms can be provided to the receiver for reception of a packet at any given time.

The design of a packet preamble and the corresponding receiver system suitable for achieving rapid synchronization to a received packet is a function of the particular spread spectrum waveform structure being used. Significant design differences exist depending upon whether PN, FH, or other waveforms are used, upon whether or not the waveform used is a function of time and, in the case of a fixed waveform, on its repetition period. A typical design is discussed in Section VII which has been implemented in an experimental packet radio network.

1) Unsynchronized Network: In designs where a fixed PN or FH spread spectrum pattern is reused for each packet, packet radio networks require no common time of day synchronization. If the repetition period of the fixed pattern is

chosen correctly, all of the benefits of a spread spectrum waveform may be realized, with the exception of the antijamming properties needed for military applications. The pattern period should exceed the maximum multipath spread expected on the channel in order to suppress intersymbol interference. For example, the PN waveform used to estimate the multipath channel depicted in Fig. 3 had a period three times the multipath spread, and allowed the channel multipath profile to be estimated without overlap. It is desirable that the pattern not repeat during the duration of a maximum length packet, so that overlapping packets interfere with each other to the minimum extent. In addition, a strong capture property is thereby realized. For ease of packet acquisition and synchronization, the fixed pattern should always be restarted from the same point at the beginning of each packet.

2) *Code Slotted Network*: In tactical applications, the waveform pattern must not be reused if the system antijamming properties are to be retained. To coordinate the waveform changing process, network time is defined in terms of discrete intervals called slots. Each slot is uniquely identified by a number. Slots are numbered sequentially, beginning at some specified time. Slot number may be viewed as the most significant digits of the system time of day. During each slot a spread spectrum pattern is defined which is unique to that slot and is used both to generate the waveform for any packet transmission beginning in that slot and as a reference input to the programmable matched filter receiver for incoming packets.

Since packet radios may operate in a random access mode, several radios could transmit overlapping packets in a given slot with the same pattern. If the slot duration is short enough, the probability of overlapping packets with identical patterns is small. This advantage from the self interference and antijamming points of view is obtained at the expense of increased difficulty of maintaining system synchronization with very small slots. Typically, one might choose a slot duration which is larger than the average direct path propagation delay by one or two orders of magnitude to facilitate synchronization and to minimize packet loss due to disagreement between transmitter and receiver as to the current slot number and corresponding valid waveform. For a typical ground radio application, this implies a slot length of a few milliseconds.

Having made a correspondence between slot number and spread spectrum pattern, a choice still remains as to the points in time at which packet transmissions may begin. Two modes, termed slotted and nonslotted, are of interest. In a slotted system, transmission is allowed to begin only at a specified point in each slot, and a guard time must be provided in the slot to allow the packet to propagate beyond normal radio range prior to the end of the slot. In this mode, the packet duration is somewhat less than a slot. Individual users make the decision independently to transmit in a given slot in random access fashion. This independence of the access contention process from slot to slot has advantageous consequences in the carrier sense mechanism described later in this section.

In contrast to a slotted system, a nonslotted system allows users to transmit at any point within a slot, and the duration of a packet transmission may exceed the duration of a single slot. The nonslotted mode may reduce system delay since the radio does not have to wait for a given time within the slot to transmit. We conjecture that this may also provide increased efficiency compared to slotted operation due to the presence

of capture. We note that this conjecture is the opposite of what is known for nonspread systems without capture. In particular, the ALOHA and slotted ALOHA [19] techniques have channel efficiencies of $1/2e$ and $1/e$, respectively. These efficiencies, however, are based on the assumption, called the zero capture model, that no packets are received correctly if they are overlapped in time at a receiver, which is not the case for spread spectrum waveform reception.

B. Capture

By the term capture, we mean the ability of a receiver to successfully receive a packet (with nonzero probability) even though part or all of the packet arrives at the receiver overlapped in time by other packets. The basic mechanism for capture is the ability of the receiver to synchronize and lock on to one packet and subsequently reject other overlapping packets as noise. A system with perfect capture is one in which the first arriving packet at an idle receiver is captured with probability one, even if another packet arrives after a vanishingly small delay.

A system cannot achieve perfect capture for a number of reasons. If the preamble portions of packets overlap, the AGC and packet synchronization process is more likely to fail, causing loss of one or both packets. If the relative time of arrival of two packets using the same pattern is less than one chip in a PN system or less than one frequency hop per interval in an FH system, successful reception of any one of the packets is unlikely. In general, there is a vulnerable period at the beginning of a packet, denoted T_c and called the capture interval, during which collision with the same portion of another packet results in the loss of both. The magnitude of T_c varies both with the waveform structure and the complexity of the receiver synchronization circuitry which can be tolerated in a given implementation. Minimizing T_c increases the capture probability.

A second major reason for imperfect capture is a result of relative received signal strength. If a packet is captured, a later arriving packet having sufficiently greater signal strength overlaps this packet, the captured packet may be destroyed. The stronger packet will in effect jam the earlier packet being received by overwhelming the interference rejection capability of the system. The likelihood that such unintentional jamming occurs depends upon the distribution of users in the network area and the distribution of signal strength rates observed at a receiver. With a typical radio implementation and assuming a uniform distribution of users in an area, the probability of such unintentional jamming can be quite low. One approach to further reducing this self-interference effect is to use transmit power control. By using cooperative network protocols, each radio can measure the signal strengths of received packets and send this information back to the transmitting radio. Using this feedback, received signal strengths may be roughly normalized by adjusting transmit power and system self-interference may be reduced. The control problems associated with the conflicting requirements of this protocol appear to be complex.

The capture phenomenon exhibits different characteristics in slotted and nonslotted systems. In slotted systems, the capture interval of the packet T_c arrives in some period of duration T_u near the beginning of the slot. If all users transmit at the same point in the slot, packets arrive in order of increasing range from the receivers. This leads to discrimination

the opposite of capture. [19] techniques effectively. These ion, called theived correctly which is not the

against the more distant users. Packet radios may implement a randomizing function which tends to distribute arrivals uniformly over the uncertainty interval T_u . If this is done, the discrimination is decreased, and capture performance improves in proportion to T_u/T_c . Unfortunately, a portion of the slot, T_u , has been lost to the communication function as a result. This loss, however, may be more than compensated by the increase in capture performance.

In a nonslotted system, T_u is equal to the slot duration. Thus, T_c/T_u is minimum and the capture probability would be expected to increase. Offsetting this factor is the potential for a packet in a nonslotted system to be interfered with by packets from adjoining slots as well as from the slot in which transmission began. It is not clear at this point whether capture performs best in a slotted or a nonslotted system, or under what conditions, and this is expected to be a subject of both theoretical and experimental study.

Receiver Addressed Waveforms

In the code slotted system described earlier, there is a requirement for a waveform generator which produces a different code spectrum pattern for each slot. If one views the mechanization of this requirement as a device which has the slot number as an input and the pattern as an output, it is easy to generalize the concept to include several input parameters and an output which is unique for each possible set of inputs. An important potential use of such a mechanism is to use the unique identifier of the next intended receiver as one input, generating a waveform which is associated with a particular slot and a particular receiver. While in the receive mode, each radio would use its own unique identifier to generate the reference pattern for its receiver. One benefit of this procedure is that contention at a given receiver is reduced to only those radios trying to send a packet to that particular receiver in the same slot. Another benefit is that the processing overhead (even the incremental routing overhead) can be reduced since only those packets intended for a given receiver are actually received by it. In networks using point-to-point routing, increased network throughput may result.

A number of difficulties arise, however, if receiver-addressed waveforms are used in place of a common code by all radios. Acknowledgment techniques must be replaced by separate acknowledgments if a hop-by-hop acknowledgment protocol is used. Network synchronization may require more overhead transmissions, since a simple broadcast on a common code pattern no longer is received by all PR's in radio range. The discovery of new RF connectivity also becomes more difficult without a common code pattern. Potential solutions to these problems include the use of dual receive channels to operate simultaneously with both a receiver-addressed pattern and a common code pattern or a time shared partition of slots into receiver addressed portion and a common code portion. The first solution requires additional hardware, while the second solution may reduce the efficiency of the shared radio channel.

Carrier Sense

The carrier sense multiple access (CSMA) technique provides a way to significantly reduce the number of conflicts due to interference on the channel. It achieves this capability by inhibiting transmissions which would otherwise cause contention to occur. The capture effect described earlier provides an alternative method to deal with contention. In a way, CSMA captures the channel, while spread spectrum allows the waveform to be captured at a receiver. Carrier sense is, in reality, a misnomer in the spread spectrum context. A better term might be "spread spectrum packet acquisition," or SSPA. However, we will use the term carrier sense to refer to SSPA for compatibility with terminology commonly used in the literature.

In carrier sense operation, a packet radio operates in the receive mode until the last possible instant prior to the onset of a packet transmission. If channel activity is present, the transmission is inhibited and rescheduled for later transmission. Perfect channel sensing is impossible in actual application due to finite propagation delay. Nevertheless, carrier sense can provide a very efficient access mechanism [11, 29].

Early models of carrier sense assumed that packet radios would implement the sensing function by measuring in-band RF energy or by observing an in-lock status indicator associated with packet synchronization circuitry. In high interference environments, however, sensing RF energy may result in false activity sensing and unnecessary transmission delays. In tactical environments, the observation of bit synchronization status may also result in false sensing if the same waveform is continually reused in the system, since this would allow a jamming transmitter to hold off the other packet radios in an area. In code slotted systems, these false sensing problems can be avoided, since new waveforms are used in each slot. A number of remaining issues related to CSMA in a spread spectrum system are addressed below.

1) *Carrier Sense Versus Capture*: Even though a packet radio system operates with a strong capture mechanism, it may still be desirable to use carrier sense to limit contention for a number of reasons. In a slotted system, sensing a packet on the channel often indicates that the potential recipients of a packet to be transmitted are committed to receiving the packet which is already on the channel. This is particularly true if receiver addressed waveforms are used, and sensing is performed using the waveform of the intended receiver as a reference.

While the capture mechanism provided by a spread spectrum waveform reduces the adverse effects of contending packets, contention may still have undesirable effects. Battery operated packet radios may waste energy unnecessarily if they transmit when it would have been possible to sense that the intended receiver was probably busy and delay the transmission. Transmitting densely overlapped packets may not result in zero capture, but the code division multiple access and anti-jamming performance of the system will be reduced. Thus, even if a strong capture mechanism exists, the use of carrier sensing procedures may allow the original margin against interference to be retained. In a slotted system the presence of carrier sense will normally increase system efficiency as well.

2) *Sensing Strategy*: Since the spread spectrum patterns in a code slotted system are constantly changing, carrier sensing can only take place practically during the preamble portion of the packet. Two alternatives are apparent. A packet radio may detect preambles and continue to receive a packet once detected. Thus, collisions with that particular packet are avoided, but any packet which arrives at the packet radio subsequent to capture of the first packet cannot be detected since the preamble will be missed. After the end of the first

packet, the channel will probably still be occupied with the end of the second packet, but, unfortunately, this will be invisible to the radio. Alternatively, if the packet radio is in a sensing mode with all local packet buffers full with packets waiting to be transmitted but with no outstanding acknowledgments, the packet radio may choose simply to sense preambles only. If a preamble is detected, transmission is prohibited for some period of time and the remainder of the detected packet is intentionally dropped. The receiver then returns to the preamble sensing mode once again in search of an idle channel interval. If during this time another preamble is sensed, the transmission is further prohibited. The difficulty with this procedure is that the location of the end of a sensed packet is unknown, due to a variable packet length. Thus, the minimum delay prior to re-attempting transmission may have to be set for a maximum length packet, adding unnecessarily to packet delivery delay.

The sensing strategies discussed above can be applied directly to the cases of slotted and nonslotted operation. However, the results of the sensing strategy are more precise in slotted operation, since packet transmission and reception must end before a slot boundary. The slot size would generally be selected to accommodate the maximum packet size plus an allowance for propagation delay and the uncertainty interval T_u if used. Thus, the slot is normally capable of accommodating only one packet per slot for a given receiver, and the sensing process may be restarted at the beginning of each slot without danger of missing a preamble due to overlapping prior transmissions.

E. Position Location

The wide bandwidth of a PN spread spectrum waveform or a hybrid PN/FH waveform allows good time resolution for measurement of packet time of arrival (TOA). In a code slotted system, a certain degree of network synchronization must be maintained to support the communication function. This degree of synchronization can be further refined to allow multiple TOA or pseudoranging measurements to be made very accurately at each packet radio with respect to other radios, and these measurements can be used in a multilateration position location procedure. Periodically, each packet radio could broadcast its current estimate of its location in order to provide adjacent units with additional data to improve the pseudoranging procedure. For typical ground-based mobile terminals, these position messages would only need to be broadcast a few times per minute, so that the load on the channel due to these messages would be a small fraction of total capacity.

In addition to the obvious applications to mobile user navigation or position reporting, knowledge of geographic relative positions and velocity vectors of packet radios could aid significantly the process of point-to-point route establishment, mobile terminal route evaluation, and other network procedures. However, the value of this tool has yet to be established.

F. ECCM Performance

A code slotted packet radio network can be designed to provide an effective electronic counter countermeasure (ECCM) capability. If the valid waveform for any given slot cannot be predicted by a network nonparticipant, jammers are constrained to operate as noncoherent rather than coherent noise generators. Jammers are also unable to "spoof" the network

with false traffic if previously used waveforms are rapidly discarded in favor of new waveforms valid in the very next slot. The decreased spectral density of spread spectrum signals, lack of waveform predictability, use of multiple access and random access, and co-located CDMA networks all contribute to a decreased susceptibility to timely detection, direction finding, and intercept of signals by nonnetwork participants. All of these properties could be important in tactical applications of packet radio.

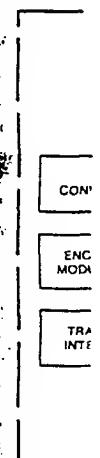
VII. EXPERIMENTAL PACKET RADIO DEVELOPMENTS

Although significant analytical progress has been achieved, a complete mathematical analysis of a multiple access, multihop radio network is not yet possible. Computer simulations having reasonable detail are being developed, and show great promise, but their practicality and degree of realism have yet to be validated. Thus, experimentation with a real network is still a primary tool of the network architect.

Toward this end, the Advanced Research Projects Agency (ARPA) initiated in 1973 a theoretical and experimental packet radio program. The initial ARPA program objective was to develop a geographically distributed network consisting of an array of packet radios managed by one or more mini-computer based "stations," and to experimentally evaluate the performance of the system. The first packet radios were delivered to the San Francisco Bay area in mid-1975 for initial testing. A quasi-operational network capability was established for the first time in September 1976, shortly after the prototype station software was developed. Approximately 25 radios are currently available for use and about 50 radios are expected to be available by the end of 1979. As of this writing (June 1978), the packet radio network has been in daily operation for experimental purposes for almost two years. During that period nearly three dozen major demonstrations of the network were scheduled and successfully carried out. The location of the major elements of the packet radio tested during 1977 is shown in Fig. 12.

The packet radio equipment currently in use in the Bay Area testbed is designed to support the development and evaluation of fundamental network concepts and techniques. This initial radio equipment was designated the Experimental Packet Radio (EPR), and is briefly described below. A major new development in 1978 has been the completion of an Upgraded Packet Radio (UPR) which is similar in architecture to the EPR but which has, in addition, the ECCM features necessary to verify the viability of packet radio concepts in tactical military applications. Selected features of the EPR and UPR design are described in this section.

Functions provided within the station software installed in 1977 included: network routing control; a gateway to other networks; a network measurement facility which collects, stores, and delivers experimental statistics from any network components; a debugging facility which supports examining and depositing the contents of memory in the PR units; an information service which assists in locating and connecting people currently using the PRNET; and an experiment configuration control module. Of these functions only the network routing control must reside within the station. All the other functions can reside in separate hosts attached to the PRNET. This permits convenient expansion of the quantity and quality of services provided to users of the PRNET. At the same time the station software itself may be kept small and simple, permitting economical replication for high redundancy and



ability of net
major ob
development an
operational
selected com
have bee
performanc

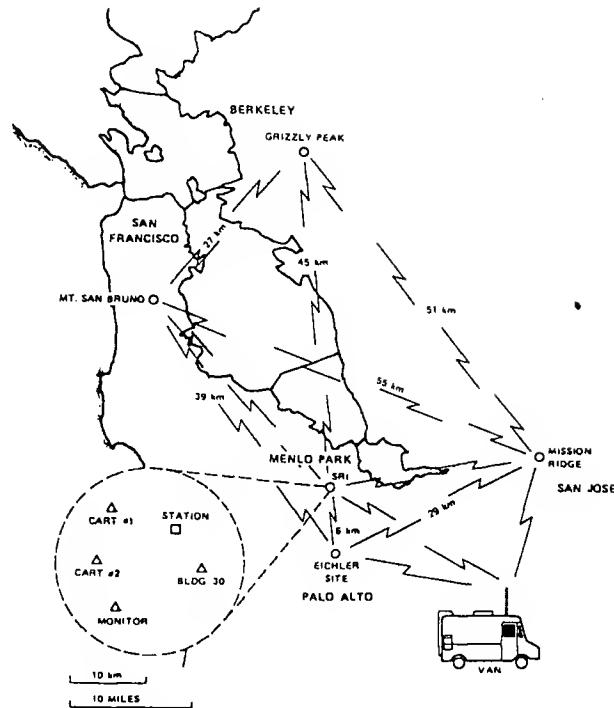


Fig. 12. Location of major elements of the packet radio testbed during 1977.

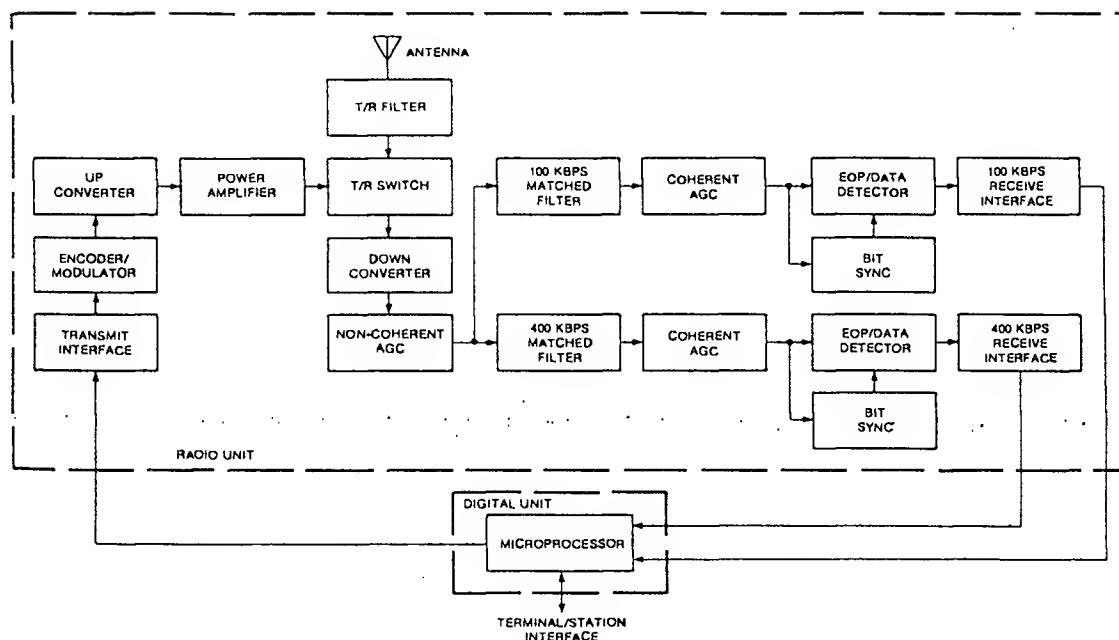


Fig. 13. Experimental packet radio configuration.

ability of network operation.

All the objectives of the experimental testbed are the development and evaluation of efficient network protocols to the PRNET, the quality and quantity of operational techniques and the test use of these capabilities.

At the present time, selected communities of real users. Toward this end, techniques have been developed to allow a variety of debugging and performance measurement operations to be carried out

in the network, and these capabilities are also described in this section.

A. An Experimental Packet Radio Unit

The EPR design which has been implemented for use in the testbed is functionally configured as in Fig. 13 (also see Fig. 6). An in-depth discussion of the implementation issues

and applied technology relevant to the EPR is available [26]. The description provided here focuses on the operation and evolution of the EPR in the packet radio test bed.

Each EPR consists of a radio unit, which transmits and receives packets, and a microprocessor-based digital unit, which controls the radio and provides packet header processing (e.g., for routing of packets between nodes). An EPR may operate as a repeater, or may be connected to a user's host computer or terminal, or to a station. The interface between the user equipment and the EPR digital unit is the portal through which packets enter and leave the network.

The EPR radio unit operates with a fixed PN spread spectrum pattern which, for simplicity in implementation, is identical for each transmitted bit. Two transmission data rates are available, 100 and 400 kbytes/s, with corresponding spread spectrum patterns of 128 and 32 chips per bit, respectively. The 100-kbytes/s rate is used for links with potentially large multipath spreads because the fixed bit length PN chip pattern does not provide the ability to discriminate against intersymbol interference. The radio unit operates in a half duplex mode. When a packet is transmitted, the preamble, header and text are read from microprocessor memory under direct memory access (DMA) control. The radio unit completes the packet format previously illustrated in Fig. 8 by adding a 32 bit cyclic redundancy checksum (CRC), then differentially encodes the data, and adds (modulo two) the appropriate PN chip pattern for the selected data rate. The resulting PN modulated stream is then applied to a minimum shift keying (MSK) modulator, and the signal is up-converted to a selected 20 MHz portion of the 1710-1850 MHz band, power amplified, and transmitted through an azimuthally omnidirectional antenna.

When not transmitting, the EPR remains in the receive mode. An arriving packet proceeds through RF amplification, down-conversion, IF amplifier and wide-band (noncoherent) automatic gain control (AGC) functions. Because the PN chip patterns used for the 100 and 400-kbits/s data rates are chosen for low cross correlation performance, two parallel receive chains following the IF amplifier/AGC can be simultaneously active. A fixed surface acoustic wave (SAW) device is used to match filter the PN modulated waveform two bits at a time in a differential detector. The outputs of the differential detector (i.e., the correlation pulses) are used to set a narrowband (coherent) AGC, and after a phase locked loop bit sync circuit has settled and the end of preamble (EOP) bit pattern has been detected, the differentially detected data in the active channel is passed to microprocessor memory under DMA control. The packet preamble design illustrated in Fig. 7 allows AGC settling, bit sync acquisition, and EOP detection (data sync) to occur prior to arrival of the first data bit. The microprocessor executes the appropriate protocol software to determine whether the received packet should be relayed, delivered to an attached user or station, or discarded.

The EPR digital unit presently uses a National IMP-16 micro-processor with 4096 16-bit words of RAM and 1024 words of PROM. The PROM contains the processor operating system and all DMA I/O routines. Network protocols and packet buffers are stored in RAM. Eventually, the protocol routines will also be stored in PROM, but the evolving nature of the software in an experimental system requires the use of RAM at present. Protocols currently implemented are the channel access protocol (CAP), the reliable station to PR protocol (SPP), a statistics gathering feature called CUMSTATS, and a debugging package called X-RAY. CAP is responsible for the

primary EPR function of transferring packets to or from adjacent EPR on a route through the network. CAP is responsible for monitoring the hop-by-hop echo acknowledgement process, retransmission of nonacknowledged packets, invoking alternate routing procedures, and determining packet disposition. CAP currently implements pure ALOHA, CSMA, and a variant of pure ALOHA in which random transmission is deferred until the end of an on-going reception process so as to needlessly discard an arriving packet.

SPP is an end-to-end protocol which is used for reliable delivery of network monitor and control packets, such as "labelled packets" sent to the EPR. SPP uses CAP as its interface to the network. Similarly, an external user device which submits a packet to the EPR uses CAP to launch the packet into the network, but is responsible for any user end-to-end protocol which might be required to be layered on top of CAP. Currently, approximately 1000 words of memory are devoted to the EPR operating system, 600 words to the DMA I/O routines, 2500 words to CAP, SPP, and CUMSTATS, and 900 words to packet buffers.

B. An Upgraded Packet Radio Unit

1) Motivation: The EPR was designed to allow experimentation and verification of the majority of basic packet radio protocols and concepts. The need exists, however, for a number of packet radios having enhanced capabilities in order to demonstrate the feasibility and advantages of packet switched radio networks in ground-mobile tactical environments. The upgraded packet radio (UPR) design which has been developed in response to this need differs from the EPR design primarily in its enhanced electronic counter measures (ECCM), a probability of intercept and position location capabilities.

In the tactical environment, reuse of the same PN sequence would result in vulnerability to jamming and waveform spoofing. In the UPR, a PN pattern, which varies on a bit-by-bit basis, is used to spread spectrum modulate each bit. As a result, the UPR must have a programmable matched filter to receive the PN modulated waveform, and the UPR network must maintain a degree of synchronization among network elements to enable the receiver to generate the same PN sequence as the sender. Timing is provided by means of an accurate time of day clock maintained by all the UPR's. In addition to varying the PN pattern used, a higher spread factor (or number of chips per bit) is used in the UPR than in the EPR. The data rate in the UPR is approximately the same as in the EPR, but the corresponding bandwidth of the UPR is larger, approximately 140 MHz (i.e., 1710-1850 MHz). An electroacoustic convolver is used in the UPR to provide the wideband programmable matched filtering element. This convolver is a single passive device which currently occupies a few cubic inches and is equivalent to many conventional matched filters. Further discussion of the operation of the convolver is contained in reference [30], [31].

Reliable packet transport in the EPR system relies on error detection and retransmission strategies. These techniques are inadequate in a tactical environment because they allow simple jamming strategies to force large number of packets retransmissions with a low attendant throughput. The U79 is provided with a forward error correction (FEC) mechanism based on convolutional encoding and sequential decoding which operates in combination with error detection and packet retransmission techniques.

2) Organization and Basic Operation of the UPR: The ~~UPR~~

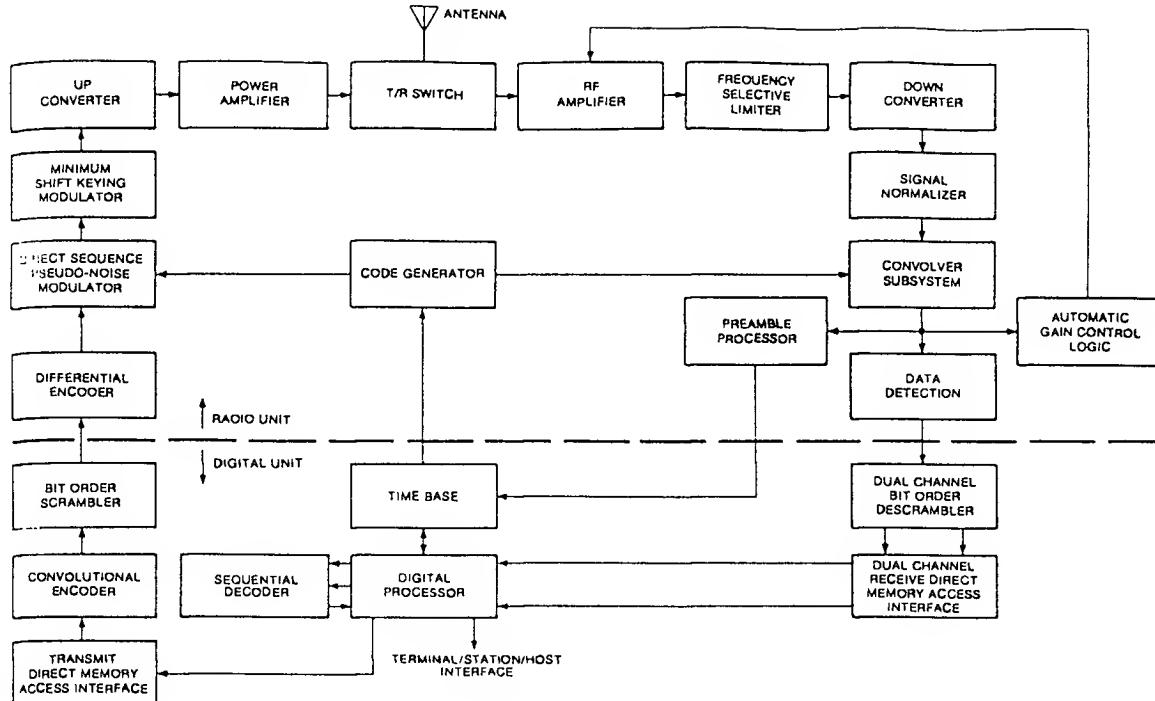


Fig. 14. Upgraded packet radio configuration.

onal organization of the UPR is illustrated in Fig. 14. As in the EPR, the UPR digital unit interfaces with user terminal equipment by means of a direct memory access (DMA) channel and with the radio unit by means of DMA channels under processor control. In addition, the UPR processor controls the local time of day clock through addressable registers and supervises the activity of the sequential decoder which operates read and write of the processor memory by means of three DMA channels.

Network synchronization is achieved through the use of time slotted packets. Transmitted packets contain their time of transmission as determined from the local time of day clock. When packets are received, the time of arrival is stored in memory. The processor uses this data to develop timing corrections for the local time base in accordance with network timing protocols.

System time is divided into basic units called slots which have 1 ms duration. For experimental purposes, however, the slot duration is adjustable. Slots are uniquely identified by sequentially numbering them. In order for two network UPR elements to communicate, their time of day clocks must agree on the current slot number. There is a one-to-one correspondence between a slot number and a PN code sequence used to transmit a packet such that when a UPR initiates a packet transmission in a particular slot, the PN code sequence used is the one corresponding to that slot. The transmission is allowed to begin at any point in the slot and the duration of the transmission may be many slots. The same strategy is used at a UPR which is waiting to receive a packet. At the beginning of each time slot, the time base delivers a new slot number to the PN code generator. This slot number completely determines the PN pattern used to receive any packet whose reception begins during that slot. Regardless of the time of arrival of the

packet within the slot, or its length, the PN code sequence used is completely defined by the slot in which reception begins. Thus, the UPR implements a code slotted system with nonslotted transmission. The hardware is also capable of slotted transmission for experimental purposes.

When a packet is to be transmitted, the processor activates a DMA channel to control and monitor the transmission. Under DMA control, the packet is read from the processor memory, convolutionally encoded with a constraint length 24 code, and loaded into a buffer prior to scrambling (bit order permutation). The packet data is read from the buffer bit by bit in pseudo-random order, differentially encoded, and passed to the spread spectrum modulator where each data bit is modulo two added to each chip of the PN chip sequence used to encode that bit. The PN modulated chip sequence is then passed to an MSK modulator, implemented with a SAW device, and having an IF output at 300 MHz. This signal is up-converted to 1780 MHz, amplified to 10 W, and fed to the azimuthally omnidirectional antenna. Fig. 15 shows the basic UPR packet and preamble format. In the discussion above, only the header and text bits of the packet are read from the processor memory. The preamble and postamble bits are supplied by the code generator circuitry, and are used in combination by the receiver to determine the receive data rate and coding format of the packet.

The UPR operates in a half duplex mode. When the UPR is not transmitting, received packets pass from the antenna through a number of RF amplifier/automatic gain control (AGC) stages. The signal is then processed by two frequency selective limiter (FSL) stages which provide an adaptive notching mechanism for narrow-band interference. After down conversion to 300 MHz, the signal is processed by a signal normalizer which tries to normalize wideband interference to the same level as the desired signal prior to spread spectrum

processing. The normalizer is a signal processing device which operates at 1F. When used in conjunction with the spread spectrum processing gain, it can nullify the effects of a single wideband source of interference. The UPR can also be operated with an adaptive antenna array to null wideband jammers.

The convolver subsystem provides the spread spectrum matched filtering function. During the preamble of the packet, the convolver outputs are further processed by a preamble processor where packet detection and synchronization functions are performed. During the remainder of the packet, the convolver outputs are processed by the data detection circuitry which makes hard binary bit decisions and provides a matched filter function for the postamble sequence. Received packets are buffered in one of the receiver's descramblers prior to bit reordering and storage of the packet in processor memory under control of a DMA channel. Two receive descrambler/DMA channels are provided to allow reception of two successive packets with minimum interpacket arrival time. When a packet has been received and stored in memory, the processor initializes DMA transactions with the sequential decoder to decode the packet. The decoding process thus takes place "off-line" in that additional packets may be received in memory while previous packets are still being decoded. As soon as the packet header is decoded, it may be submitted to network protocol processing to determine disposition of the packet.

3) *Spread Spectrum Code Selection and Generation:* One of the fundamental differences between the EPR and UPR equipment is the use in the UPR of a direct sequence spread spectrum waveform which changes from bit to bit. The selection of the code and an efficient means for its generation was an important consideration in the UPR development. The attributes of the code which influenced the choice were as follows. In order to provide good antijamming and antispoofing properties, it should be impossible to predict the PN sequence which will be used to encode some future data bit based on an observation of the PN code used for encoding past data. The PN code sequences of bit duration should also have good autocorrelation properties to facilitate detection and synchronization of the packet. Code autocorrelation becomes even more important in view of the position location subsystem being developed for the UPR which relies on accurate time of arrival measurement for its operation. The PN code sequences used to encode bits should have low cross-correlation. This attribute is important since it provides a capture mechanism for the channel access technique and will support multiple simultaneous packet transmissions using code division multiple access (CDMA). The equipment required to generate the code should require low power and occupy low volume. The PN chip stream rate is quite high, and low power is therefore of particular concern. In the current application, it is desirable to have a code generation algorithm which is easily time reversible, since the convolver reference chip stream is the reverse of the modulator chip stream over a bit. Any code generator can be made reversible by running it in advance, storing the chips in a memory, and reading them in reverse order. However, a reversible code algorithm may result in savings of power and hardware.

A code generation technique having the desired attributes was selected from several candidates for implementation in the UPR. For brevity, a code pattern will refer to the length of PN chip sequence used to spread spectrum encode one data bit. In the initial versions of the UPR, a finite set of approxi-

mately 2000 code patterns are being used. For any pattern in this set, its chip-by-chip binary complement is also in the set. This feature makes it more difficult for an unauthorized user to extract the data from a packet. For any given transmitted packet, a sequence of codes is constructed from this finite code set according to a nonlinear secure algorithm. The sequence constructed depends on the slot number during which transmission begins and certain other parameters used by the algorithm. Use of this technique allows a good algorithm operating at low power and low speed to specify a high speed chip stream having the desired unpredictability to an observer outside of the system. In addition, the autocorrelation and cross-correlation properties desired are a function only of the finite set of codes selected. For the UPR equipment, the code set was constructed from the Gold codes [32], which provides the desirable property of reversibility and which can be implemented in a memory structure using parallel operations at low speed and low power consumption prior to a final high speed parallel to serial operation.

4) *Error Control:* The UPR employs convolutional encoding and sequential decoding for forward error correction in conjunction with the error detection and retransmission techniques used in the EPR [33]. As indicated in Fig. 15(a), only the header and text of the packet are encoded. The convolutional encoder is capable of encoding at three rates: 1/4, 1/2, and 3/4. Two code rates applied to each coded packet. A lower rate code is used for the first portion of the packet and a higher rate is used for the larger second portion. The length of the first portion is under software control and need not be identical to the header. The code rate combinations allowed in the header and text, respectively, are (1/4, 1/2), (1/2, 3/4), and (1/2, 1), where code rate 1 indicates that the encoder is bypassed. The UPR processor submits the two distinctly coded portions of a received packet to the sequential decoder as separate tasks. This allows the processor to decode one or more packet headers quickly so that header protocol processing may take place while the longer portions of the packet are being decoded.

If the decoding time of the second portion becomes excessive, the decoding can be aborted by the processor and the original packet can be saved in the encoded state pending arrival of a retransmission. If the first portions, which are more heavily encoded, of both the original and the duplicate can both be decoded, the two packets can be identified as duplicates. Since the encoded second portions of the two packets are identical, except for their error sequences, they can be combined to form a virtual binary erasure channel. The sequential decoder is designed to operate on either binary symmetric or binary erasure channel data. The result of this procedure is that if the first portions of the packets were decodable at, say, rate 1/4 as separate binary symmetric channel inputs, then the combined second portions will, with high probability, be successfully decoded as a rate 1/2 binary erasure channel input.

It should be pointed out that sequential decoders operate best with random channel errors. For this reason, a bit order randomization across the entire packet is performed after coding, and the inverse operation is performed prior to decoding. The randomization algorithm is a function of the slot number and packet length. As a result, a jammer cannot impose the same error pattern in two copies of a packet, and burst errors in the channel appear as random errors to the decoder. This may also significantly increase the power of the binary erasure combining technique.

indicated sync pattern the purpose as to wh applies to t bit rate su(s) and whe could be bypassed. an uncode dancy check

Processor design is c Original EPR functions results of ne system thr channel limited RF spec

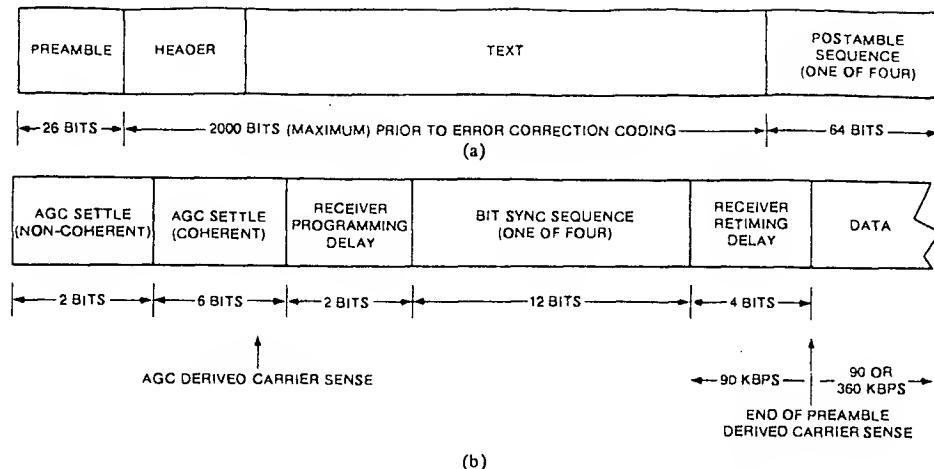


Fig. 15. (a) UPR packet format. (b) UPR packet preamble detail.

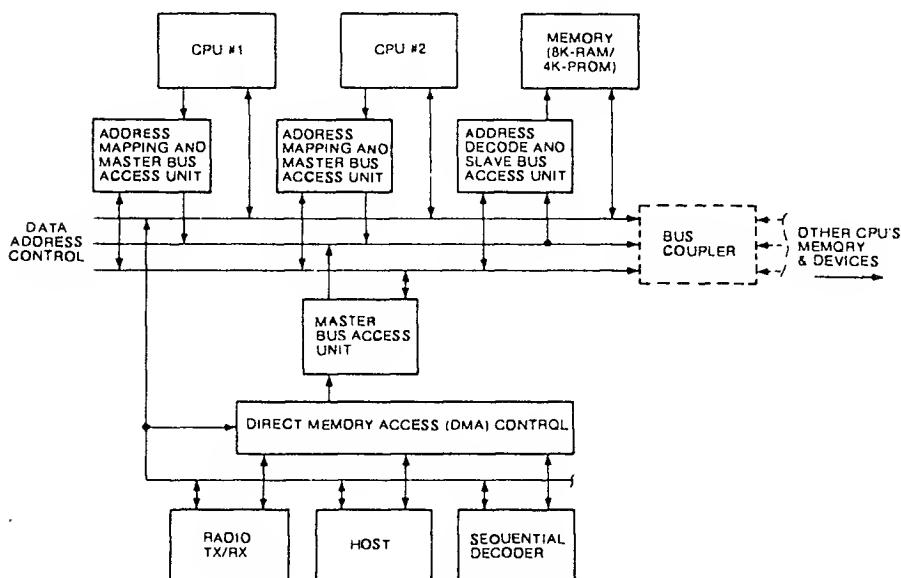


Fig. 16. UPR digital section architecture.

As indicated in Fig. 15, there are four distinct preamble sync patterns and four distinct postamble patterns. Part of the purpose of these patterns is to inform the receiving PR as to which forward error correction encoding format applies to the packet. The other information imparted is bit rate subsequent to the preamble (90 kbit/s or 360 kbit/s) and whether or not the bit order descrambling function could be bypassed. The latter feature is used in conjunction with an uncoded packet format which uses only a cyclic redundancy checksum for error detection, as was the case in the PR.

Processor Architecture: The processor being used in the PR design is considerably more powerful than that used in the original EPR. The added capability is driven both by the additional functions which the UPR processor must perform, and the results of network tests in the EPR test bed which indicate that system throughput can be processor-limited rather than channel limited. Since efficient multiple access use of the limited RF spectrum resource is a key program objective, the

UPR processor should provide sufficient processing power to allow maximum use of the channel.

Fig. 16 shows the general architecture of the UPR digital section. The central processing units (CPU's) are Texas Instruments 9900 microprocessors. These microprocessors have the capability for implementing a multiprocessor architecture by providing the appropriate monitor and control signals for bus sharing. Analysis showed that for the UPR application two CPU's could most efficiently share a bus directly, but that for larger numbers of CPU's higher net efficiency was obtained using bus couplers similar to Texas Instrument's TI-Line architecture. The UPR will initially be configured with the dual CPU processor, and will have the capability to add additional processing power using bus couplers and additional CPU and memory hardware. The dual CPU processor increases the processing power available to the UPR to support additional functions such as position location, network synchronization, and distribution of spread spectrum keys. A picture of the UPR is shown in Fig. 17.



Fig. 17. An upgraded packet radio.

The basic architecture of the multiprocessor digital unit is compatible with the EPR radio section and the latest EPR units are being constructed using the dual processor design. These units will provide the added processing power desired at high traffic levels in the EPR test bed.

6) *Mobile Operation*: An L-band radio operating from a mobile platform in a nonsited, ground environment encounters severe fading due to multipath signal components. The EPR equipment was originally designed to lock onto a single multipath component during the packet preamble and to track this component for the remainder of the packet. This was a simple and reliable approach for initial experiments which did not involve mobile platforms. When operating from a platform moving at sufficient velocity, however, fading of the acquired multipath component could occur rapidly enough that portions of the packet would be lost. In order to combat this fading phenomenon, a diversity mechanism based on post-detection integration of multipath components is provided in the UPR. The simplicity of the mechanism has allowed the EPR equipment to be retrofitted with similar circuitry, so that reliable mobile operation is currently available in the EPR test bed. A block diagram of the receiver is shown in Fig. 18.

C. Network Management and Operation

At the present time two packet radio experimental networks are operating: an experimental testbed network covering much of the San Francisco Bay Area; and a local distribution network in the Boston Area, which is used for station software development.

1) *Network Monitoring and Control*: A centralized network management facility (NMF) has been developed for managing and operating the Bay Area experimental testbed. It is somewhat similar to the ARPANET network control center (NCC) [39] in that it collects and displays relevant PRNET status information on a continuous basis. The facilities available at the NMF include two stations and a network monitoring system employing an Interdata 70 and an EPR to monitor radio traffic (during test and debugging only). From the NMF, the network can be debugged, the status of the network can be monitored, tests and measurement experiments can be run, and faults can be detected, diagnosed, and isolated.

2) *Debugging the Network*: All elements of the packet radio

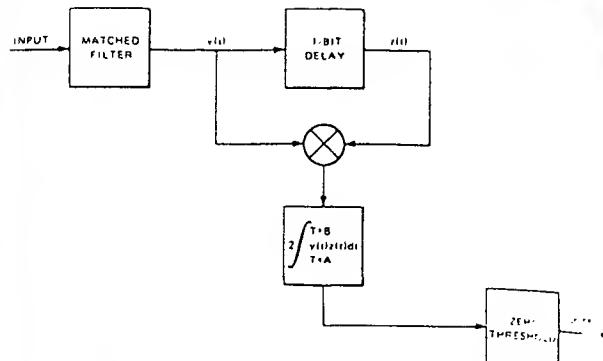


Fig. 18. DPSK receiver using post-detection integration.

network have been designed to be debugged remotely under test as well as operational conditions. The memory of an EPR packet radio's microprocessor can be remotely examined and altered through the use of the X-RAY debugger by a person at the station. The X-RAY process is routinely used to alter operating parameters in the packet radios (such as power output, frequency, timing, and protocol values) and to examine or alter program code. Its operation can also be automated so that other station processes can utilize it directly.

The PRNET is normally connected to the ARPANET [1]. This connection is accomplished using a gateway [34] programmed with the network station processor, to communicate with an ARPANET IMP [2]. The station can then be remotely debugged from an authorized ARPANET host using a cross internetwork debugger known as X-NET. By using internetwork protocols to access the station's X-RAY process, even the radio can be remotely debugged from the ARPANET. The user terminal interface unit (TIU) also supports both forms of remote debugging. This feature has proved essential to the Bay Area PRNET development in that station software developers located in Boston and packet radio software developers in Texas can remotely participate in network debugging as need arises, and new software versions can be conveniently installed from remote development sites as frequently as required.

3) *System Monitoring*: Once initialized, each packet radio in the network periodically announces its existence by transmitting to the station summary ROP's which contain neighbor tables and other status information. Similarly, terminal devices periodically send summary TOP's (terminal-on packets), which serve much the same function as their counterpart summary ROP's.

Both the station and the network monitor make extensive use of summary ROP's and TOP's. The station maintains a connectivity matrix based on the information contained in the ROP's for assigning routes. Current network connectivity can be displayed at the station upon request, and all state changes for nodes and links may be time stamped and logged. While active, the independent network monitoring system also listens to ROP's, and maintains a table of the last time that ROP's and TOP's were heard, for each packet radio or terminal interface unit ID. Thus, the exact time of failure of any network element can be obtained—even if a component of the station fails.

D. Test and Experimentation

This section describes the measurement capabilities that have been designed-in as an integral part of the experimental net-

work. Measurements, including software measurement, experiment completed, the operating software collection of data for delivery. No adjustments on through design goal should be discussion of this issue. From the station, the network is tested, if appropriate devices are off, and function of a message is followed over the analysis. The four primary cumulative packets, and neighbor activity count disposition of up packets since, and pick route to their history. Neighbors derived from these built-in TIU measure controlled traffic; collects selected data traffic; collects histogram data. The PRU measures snapshots; periodically sends PRU CUI and packets programs. Station measures the resulting process and cumstats work connection on the measurement. Off-line measurement PRNET measurement data are sent at UCLA, these facilities network experiments data collection, reduction and permanent, facilities are a future, it is likely capabilities.

work. Measurement facilities have been built into the PR and station software. They provide for the collection and delivery of measurement data over the radio channel in real-time while an experiment is being run; after the experiment has been completed, the data are reduced and analyzed at a remote site. Operating software in the PRU's, TIU's, and station performs the collection of measurement data and uses the system protocols for delivery of this data to a measurement file located at the station. No additional hardware is needed to make measurements on throughput, delay and several other parameters, and the design goal of the facilities is that the taking of measurements should have a minimal impact on PRNET performance. A discussion of PR measurement techniques is given in Tobagi *et al.* in this issue [29].

From the station, parameters in each PR and terminal device in the network can be set remotely, selected elements can be tested, if appropriate, the collection of statistical data from selected devices may be enabled, traffic sources may be turned on or off, and data collection may be initiated. At the conclusion of a measurement run, the data can be automatically loaded over the ARPANET to a remote site (e.g., UCLA) for analysis.

The four primary measurement tools that have been developed are cumulative statistics (CUMSTATS), snapshots, pickup packets, and neighbor tables. CUMSTATS consist of a variety of activity counters in each node. Snapshots periodically record the disposition of packet buffers and other node resources. Pickup packets are "crates" that start out empty at a traffic source, and pick up information at each node they traverse en route to their destination, thus providing a trace of their history. Neighbor tables are a table of counts of packets received from each "neighbor" PR in range. The location of these built-in measurement tools are as follows.

1) TIU measurement software—provides sources and sinks for controlled traffic streams; generates and collects pickup packets; collects end-device CUMSTATS; and periodically sends selected data to station measurement process. End-device CUMSTATS collected consist of packet activity counters, retransmission histograms, and end-to-end acknowledgment time delay spectra.

2) PRU measurement software—collects subnet CUMSTATS snapshots; enters local data into pickup packets; and periodically sends collected data to station measurement process. PRU CUMSTATS include counters for packets transmitted, packets received, packets in error, and retransmission histograms.

3) Station measurement software—controls experiments and collects the resulting measurement data. Internal packet-handling processes (such as the forwarder and gateway) also collect cumstats which are written on the measurement file. Network connectivity, labeling, and route updates are all written on the measurement file as they occur.

4) Off-line measurement software—the final destination of PRNET measurement data is the UCLA 360/91 computer. Data are sent from the station over the ARPANET and are stored at UCLA, for use by several analysis programs.

These facilities as described provide an advanced set of tools for network experimentation. Controlled experiments can be data collected and transmitted to a general-purpose machine for reduction and analysis, all using elements that are a built-in permanent part of the network under test. Today these facilities are an accepted part of experimental networks; in the future, it is likely that many operational networks will have similar capabilities as well.

VIII. CONCLUSIONS

Packet radio has emerged as a viable technology for both fixed and mobile computer communications. The elements of a PRNET are already small enough (the repeater is about a cubic foot) to envision the day when a complete packet radio terminal will be about the size of a pocket calculator. A packet radio is still too expensive for widespread usage outside the military. However, the application of recent advances in integrated circuit technology and sophisticated signal processing techniques is expected to lead to low cost radios within the next five to ten years.

The cost of digital technology and advanced signal processing is steadily decreasing. It is very reasonable to expect that digital radio technology will be married with the personal computer technology and that a very powerful personal network technology will emerge which will be capable of revolutionary impact on both the communications and computing fields.

In the space environment, the packet radio technology is a potential candidate for organizing multiple satellites into a store-and-forward network. This architecture would entail the provision of significantly more processing within each satellite than is typical of existing satellites.

Packet radio will most likely play a major role in achieving local distribution of information, particularly when the source or destination can be mobile. Numerous cost, fabrication, and operational issues are being pursued and various regulatory bodies are evaluating the possible role of packet radio in achieving efficient spectral utilization for computer communications. However, the basic technology has now been demonstrated and it seems likely that packet radio will play a significant future role in computer communications and the local distribution of information.

ACKNOWLEDGMENT

Many of our colleagues working on the packet radio program have contributed to the ideas presented in this paper. In particular, we would like to acknowledge contributions from ARPA, Rockwell International/Collins, Bolt, Beranek and Newman, Inc., SRI International, UCLA, MIT Lincoln Laboratory, Texas Instruments, and Linkabit Corporation.

REFERENCES

- [1] L. G. Roberts and B. D. Wessler, "Computer network development to achieve resource sharing," *AFIPS Proc. SJCC*, pp. 543-549, 1970.
- [2] F. E. Heart, R. E. Kahn, S. M. Ornstein, W. R. Crowther, and D. C. Walden, "The interface message processor for the ARPA computer network," *AFIPS Proc. SJCC*, pp. 551-567, 1970.
- [3] R. E. Kahn, "Resource-sharing computer communication network," *IEEE Proc.*, vol. 60, pp. 1397-1407, Nov. 1972.
- [4] L. G. Roberts, "The evolution of packet switching," this issue, pp. 1307-1313.
- [5] L. G. Roberts, "Data by the packet," *IEEE Spectrum*, Feb. 1974.
- [6] R. E. Millstein, "The national software works: A distributed processing system," *Proc. ACM Nat. Conf.*, Seattle, WA, 1977.
- [7] R. E. Kahn, "The organization of computer resource into a packet radio network," *IEEE Trans. Commun.*, vol. COM-25, no. 1, pp. 169-178, Jan. 1977.
- [8] N. Abramson, "The Aloha system—another alternative for computer communications," in *AFIPS Conf. Proc.*, vol. 37, FJCC, pp. 695-702, 1970.
- [9] R. Binder *et al.*, "ALOHA packet broadcasting—A retrospect," in *AFIPS Conf. Proc.*, vol. 44, pp. 203-215, 1975.
- [10] T. Klein, "A tactical packet radio system," in *NTC Proc.*, New Orleans, LA, Dec. 1975.
- [11] L. Kleinrock and F. Tobagi, "Packet switching in radio channels: Part I—Carrier sense multiple access modes and their throughput-delay characteristics," *IEEE Trans. Commun.*, pp. 1400-1416, Dec. 1975.
- [12] Special Issue on Spectrum Management, *IEEE Trans. Electromagn. Compat.*, vol. EMC-19, no. 3, May 1978.

[13] D. L. Nielsen, "Microwave propagation measurements for mobile digital radio applications," *EASCON-77 Proc.*, Sept. 1977, pp. 14-2A to 14-2L.

[14] R. C. Dixon, "Spread spectrum systems," New York: John Wiley and Sons, 1976.

[15] G. L. Turin, "Communication through noisy, random multipath channels," *IRE Convention Record*, vol. 4, part IV, 1956, pp. 154-166.

[16] R. Price and P. G. Green, Jr., "A communication technique for multipath channels," *IRE Proc.*, vol. 46, Mar. 1958, pp. 555-570.

[17] V. Charash, "A study of multipath reception with unknown delays," Ph.D. dissertation, Dep. Electrical Engineering and Computer Science, University of California, Berkeley, Dec. 1973.

[18] N. Abramson, "Packet switching via satellite," in *AFIPS Conf. Proc.*, 1979, pp. 696-702.

[19] L. Kleinrock and S. Lam, "Packet switching in a slotted satellite channel," in *AFIPS Conf. Proc.*, 1973, pp. 703-710.

[20] L. Kleinrock and M. Scholl, "Packet switching in radio channels: New conflict free multiple access schemes for a small number of data users," *ICC Proc.*, 1977, pp. 22.1-105 to 22.1-111.

[21] L. Kleinrock and F. Tobagi, "Carrier sense multiple access for packet switched radio channels," *ICC Proc.*, June 1974, pp. 21B-1 to 21B-7.

[22] L. Kleinrock and M. Scholl, "Packet switching in radio channels: New conflict free multiple access schemes," to appear in *IEEE Trans. Communications*.

[23] J. Capetanakis, "The multiple access broadcast channel: Protocol and capacity considerations," Ph.D. dissertation, M.I.T., Aug. 1977.

[24] J. Burchfiel, R. Tomlinson, and M. Beeler, "Functions and structure of a packet radio station," *AFIPS Conf. Proc.*, vol. 44, 1975, pp. 245-251.

[25] H. Frank, I. Gitman and R. VanSlyke, "Packet radio system—Network considerations," *AFIPS Conf. Proc.*, May 1975, pp. 217-231.

[26] S. Fralick and J. C. Garrett, "Technological considerations for packet radio networks," in *AFIPS Conf. Proc.*, pp. 233-243, 1975.

[27] F. Tobagi and L. Kleinrock, "On the analysis of buffered packet radio systems," in *Proc. 8th HICCS, Honolulu, HI*, Jan. 1976, pp. 42-45.

[28] Y. Dalai, "Broadcast protocols on packet-switched computer networks," Ph.D. dissertation, Stanford University, Digital Systems Laboratories, technical report no. 128, Apr. 1977.

[29] F. A. Tobagi, M. Gerla, R. W. Peebles, and E. Manning, "Modeling and measurement techniques in packet communication networks," this issue, pp. 1423-1447.

[30] J. Cafarella, J. Allison, W. Brown, and E. Stern, "Programmable matched filtering with acoustic-electric convolvers in spread spectrum systems," in *1975 Ultrasonics Symp. Proc.*, IEEE CAT #75, CHO 994-4s, pp. 205-208.

[31] —, "Acoustoelectrical convolvers for programmable matched filtering in spread spectrum systems," *Proc. IEEE*, vol. 64, W, 1976.

[32] R. Gold, "Optimal binary sequences for spread spectrum multiplexing," *IEEE Trans. Inform. Theory*, pp. 617-621, Oct. 1967.

[33] J. Wozencraft and L. Jacobs, *Principles of Communication Engineering*, New York: John Wiley and Sons, 1966.

[34] V. G. Cerf and P. T. Kirstein, "Issues in packet network interconnection," this issue, pp. 1386-1408.

[35] Y. Okamura et al., "Field strength measurements and its variability in VHF and UHF land mobile radio services," in *Communication Channels: Characterization and Behavior*, B. Goldberg, Ed. New York: IEEE Press, 1976.

[36] A. Longley and P. Rice, "Prediction of tropospheric radio transmission over irregular terrain: A computer method," *RAA Tech. Report ERL 79-ITS 67*, 1968.

[37] I. M. Jacobs, R. Binder, and S. V. Hoyersten, "General purpose packet satellite networks," this issue, pp. 1448-1467.

[38] V. G. Cerf and R. E. Kahn, "A protocol for packet network intercommunications," *IEEE Trans. Commun.*, vol. COM-22, pp. 637-648, May 1974.

[39] A. McKenzie, "The ARPA Network control center," in *IV. Fourth Data Commun. Symp.*, Quebec City, Canada, pp. 5-6, Oct. 1975.